



US009349201B1

(12) **United States Patent**
Gault et al.

(10) **Patent No.:** **US 9,349,201 B1**
(45) **Date of Patent:** **May 24, 2016**

(54) **COMMAND SENTINEL**

USPC 345/501
See application file for complete search history.

(75) Inventors: **Andrew Buchanan Gault**, Newport Beach, CA (US); **Rui Filipe Andrade Pereira**, Aliso Viejo, CA (US); **David Perry**, Monarch Beach, CA (US); **Brian K. F. Lau**, Huntington Beach, CA (US); **Kelvin Michael Yong**, Irvine, CA (US); **Claes Christian Rickeby**, Trabuco Canyon, CA (US); **Ryan Hamilton Breed**, Laguna Niguel, CA (US); **Eleazar T. Galano, III**, Glendale, CA (US); **Austin English**, College Station, TX (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,889,499 A 3/1999 Nally et al. 345/7
6,102,969 A 8/2000 Christianson et al. 717/146

(Continued)

FOREIGN PATENT DOCUMENTS

WO WO 03/075116 9/2003

OTHER PUBLICATIONS

Unknown, "T5 Labs Reveals Instant Gaming Platform", Article Nov. 9, 2007, p. 1-8, <http://www.edge-online.com>.

(Continued)

Primary Examiner — Xiao Wu

Assistant Examiner — Mohammad H Akhavannik

(74) *Attorney, Agent, or Firm* — Martine Penilla Group, LLP

(73) Assignee: **Sony Interactive Entertainment America LLC**, San Mateo, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 949 days.

(21) Appl. No.: **13/350,216**

(22) Filed: **Jan. 13, 2012**

Related U.S. Application Data

(63) Continuation-in-part of application No. 12/826,489, filed on Jun. 29, 2010, now Pat. No. 8,888,592, and a continuation-in-part of application No. 12/791,819, filed on Jun. 1, 2010, and a continuation-in-part of

(Continued)

(51) **Int. Cl.**
G06F 15/00 (2006.01)
G06T 11/60 (2006.01)


(52) **U.S. Cl.**
CPC **G06T 11/60** (2013.01); **G09G 2340/125** (2013.01)

(58) **Field of Classification Search**
CPC H04L 67/02; H04L 67/10; H04L 67/42; A63F 13/12

(57) **ABSTRACT**

A video server is configured to provide streaming video to players of computer games over a computing network. The video server can provided video of different games to different players simultaneously. This is accomplished by rendering several video streams in parallel using a single GPU (Graphics Processing Unit). The output of the GPU is provided to graphics processing pipelines that are each associated with a specific client/player and are dynamically allocated as needed. A client qualifier may be used to assure that only clients capable of presenting the streaming video to a player at a minimum level of quality receive the video stream. Video frames provided by the video server optionally include overlays added to the output of the GPU. These overlays can include voice data received from another game player. These overlays may be used to prevent presentation of non-allowed input controls.

16 Claims, 13 Drawing Sheets



Related U.S. Application Data

application No. 11/499,524, filed on Aug. 3, 2006, now abandoned.

- (60) Provisional application No. 61/354,699, filed on Jun. 14, 2010, provisional application No. 61/183,035, filed on Jun. 1, 2009, provisional application No. 61/183,037, filed on Jun. 1, 2009, provisional application No. 61/183,088, filed on Jun. 2, 2009, provisional application No. 61/183,546, filed on Jun. 2, 2009, provisional application No. 61/323,354, filed on Apr. 12, 2010, provisional application No. 61/345,534, filed on May 17, 2010, provisional application No. 61/488,297, filed on May 20, 2011, provisional application No. 61/558,376, filed on Nov. 10, 2011, provisional application No. 61/558,859, filed on Nov. 11, 2011.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,241,612 B1	6/2001	Heredia	463/42
6,342,892 B1	1/2002	Van Hook et al.	345/503
7,090,577 B2	8/2006	Serizawa et al.	463/6
7,096,253 B2	8/2006	Vinson et al.	709/203
7,227,894 B2	6/2007	Lin et al.	375/240.12
7,275,987 B2	10/2007	Shimakawa et al.	463/9
7,448,063 B2	11/2008	Freeman et al.	725/136
7,804,856 B2	9/2010	Krause et al.	370/486
7,818,769 B2	10/2010	Peacock et al.	725/54
7,916,147 B2	3/2011	Clemie et al.	345/522
7,953,883 B2	5/2011	Thomas et al.	709/231
8,081,192 B2	12/2011	Clemie et al.	345/522
8,095,400 B2	1/2012	Herde et al.	705/5
8,152,631 B2	4/2012	Anderson	463/25
8,203,568 B2	6/2012	Clemie et al.	345/522
2003/0083132 A1	5/2003	Berg et al.	463/40
2003/0190952 A1	10/2003	Smith et al.	463/30
2003/0228906 A1	12/2003	Walker et al.	463/40

2005/0054423 A1	3/2005	Wadleigh	463/20
2005/0104889 A1 *	5/2005	Clemie et al.	345/522
2006/0121991 A1	6/2006	Borinik et al.	463/43
2006/0146057 A1	7/2006	Blythe	345/506
2006/0218225 A1	9/2006	Hee Voon et al.	709/201
2006/0224761 A1	10/2006	Howarth et al.	709/231
2006/0230428 A1	10/2006	Craig et al.	725/133
2007/0088852 A1 *	4/2007	Levkovitz	709/246
2007/0155507 A1	7/2007	Gatto et al.	463/42
2007/0265094 A1	11/2007	Tone et al.	463/42
2008/0207322 A1	8/2008	Mizrahi	463/32
2008/0234047 A1	9/2008	Nguyen	463/42
2008/0284798 A1	11/2008	Weybrew et al.	345/630
2008/0291208 A1	11/2008	Keall	345/506
2009/0118017 A1	5/2009	Perlman et al.	463/42
2009/0276402 A1	11/2009	Stiers	707/3
2010/0167816 A1	7/2010	Perlman et al.	463/30
2010/0248817 A1	9/2010	Anderson	463/25
2010/0317443 A1	12/2010	Cook et al.	463/42
2012/0004039 A1	1/2012	Perry et al.	463/42
2012/0004040 A1	1/2012	Pereira et al.	463/42
2012/0004041 A1	1/2012	Pereira et al.	463/42
2012/0004042 A1	1/2012	Perry et al.	463/42
2012/0064976 A1	3/2012	Gault et al.	463/42
2012/0172124 A1	7/2012	Estrop et al.	463/31
2012/0200583 A1	8/2012	Clemie et al.	345/522

OTHER PUBLICATIONS

Parfitt, "Pay for Play" gaming a reality—Article Nov. 20, 2007, p. 1-7, <http://www.mcvuk.com>.
 Schramm, Onlive calls T5 Labs claim to key game streaming patent "irrelevant"—Article Feb. 15, 2011, VentureBeat.
 Takahashi, "Who invented cloud gaming? T5 Labs tangles with Onlive"—Article Feb. 15, 2011, p. 1-4, venturebeat.com.
 Onlive Fans blog—Blog Feb. 17, 2011, p. 1-8, <http://onlivefans.com>.
 Unknown, "New Virtual Game Console From T5 Labs Makes PC Games Playable on Televisions"—Article Nov. 12, 2011, PR Newswire UK, <http://www.prnewswire.co.uk>.

* cited by examiner

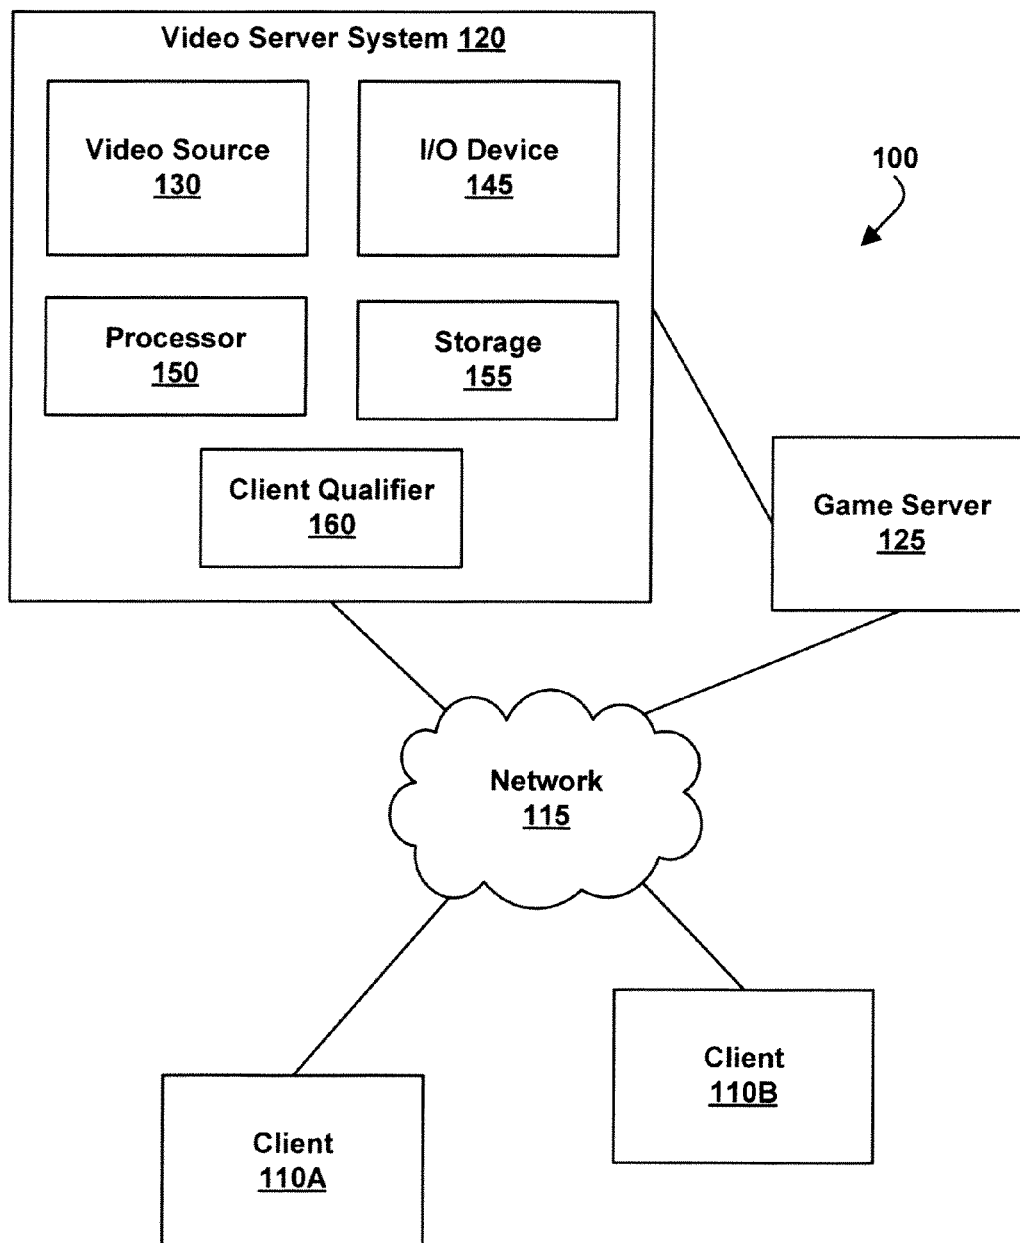


FIG. 1

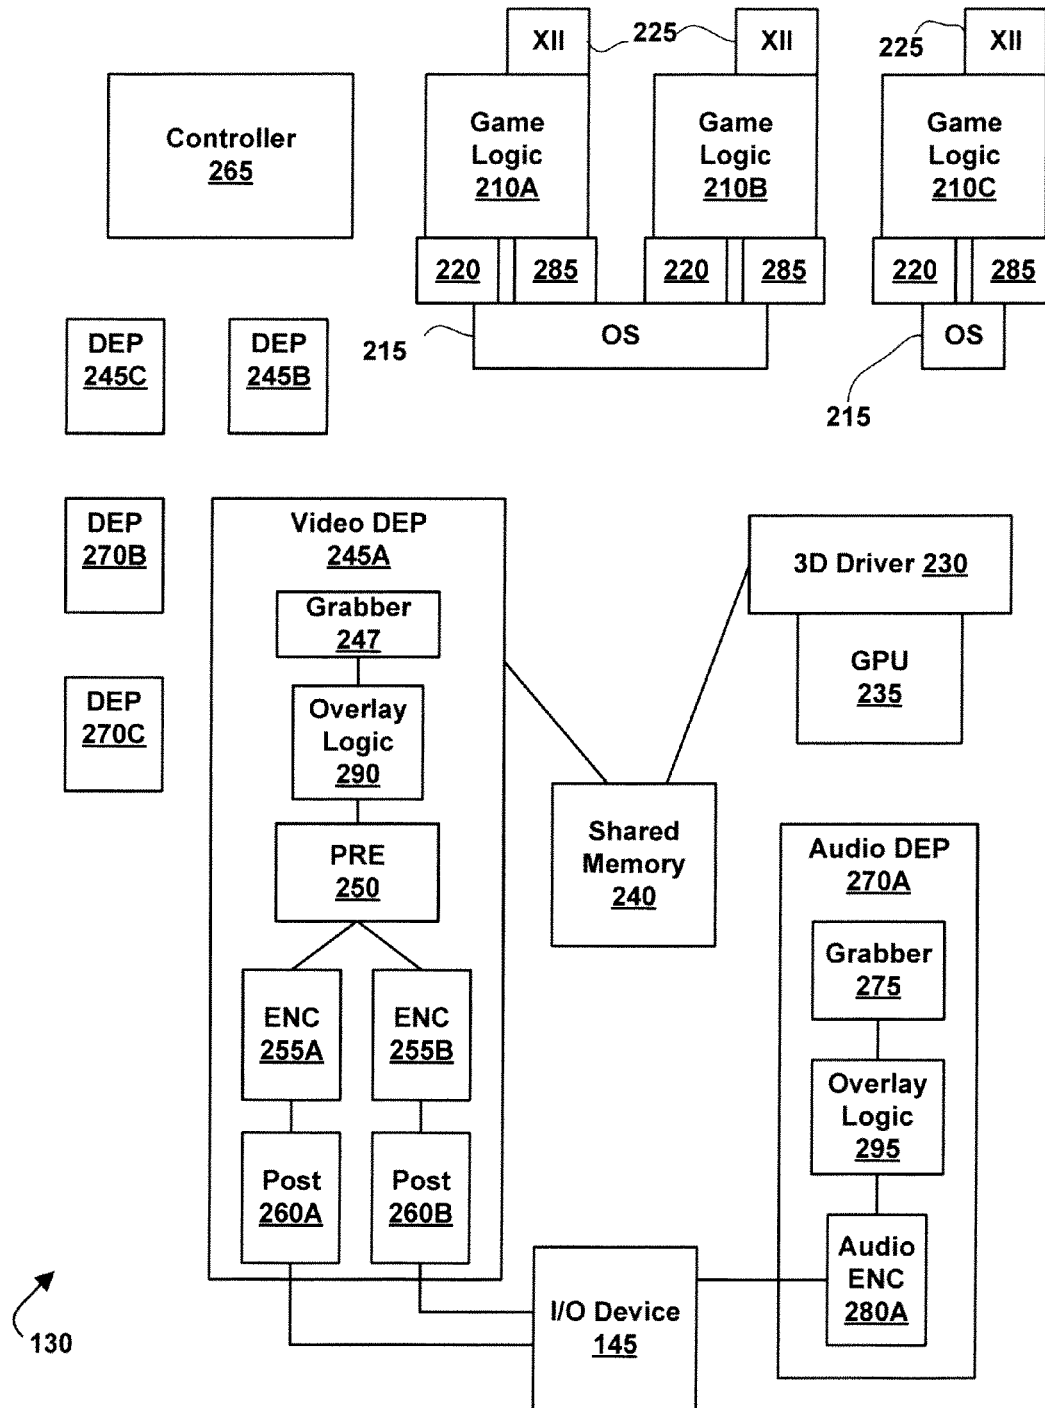
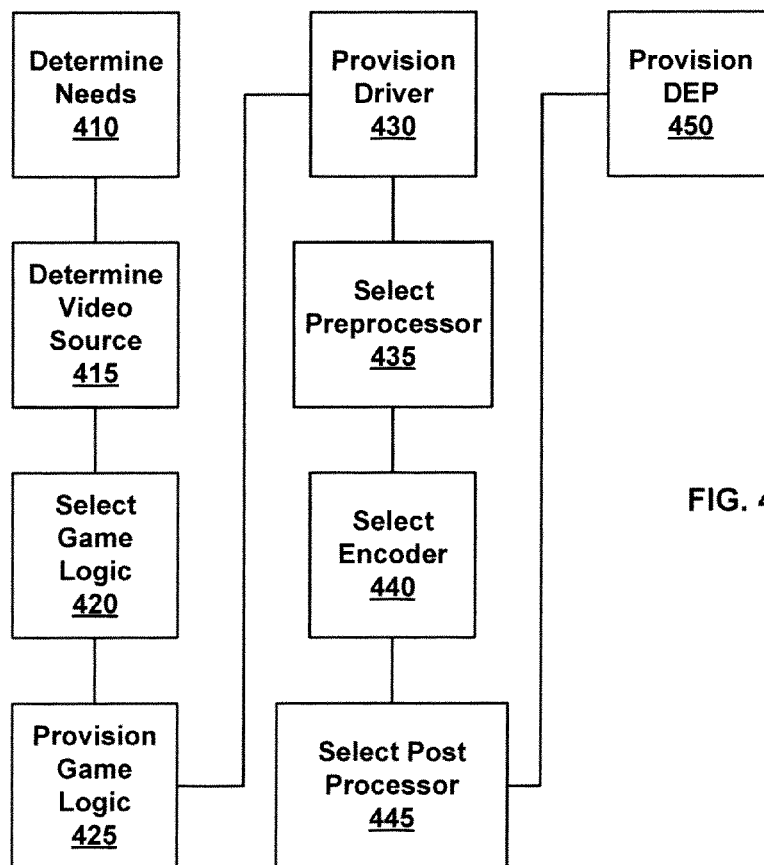
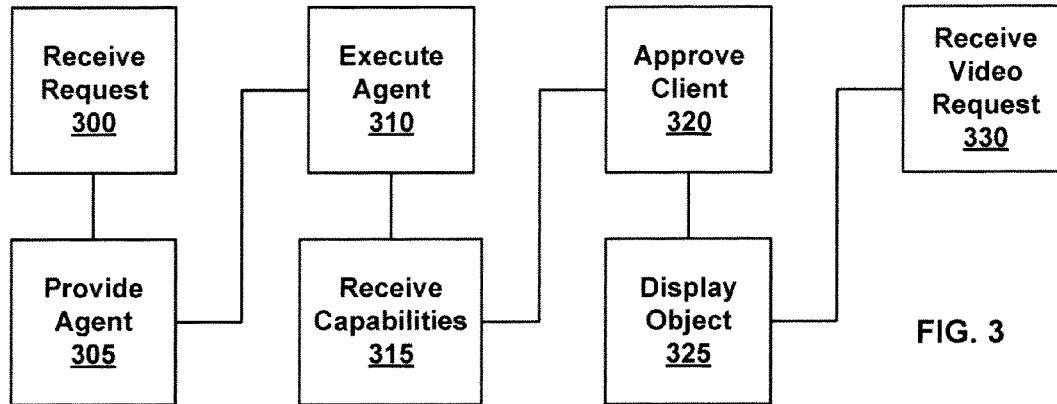
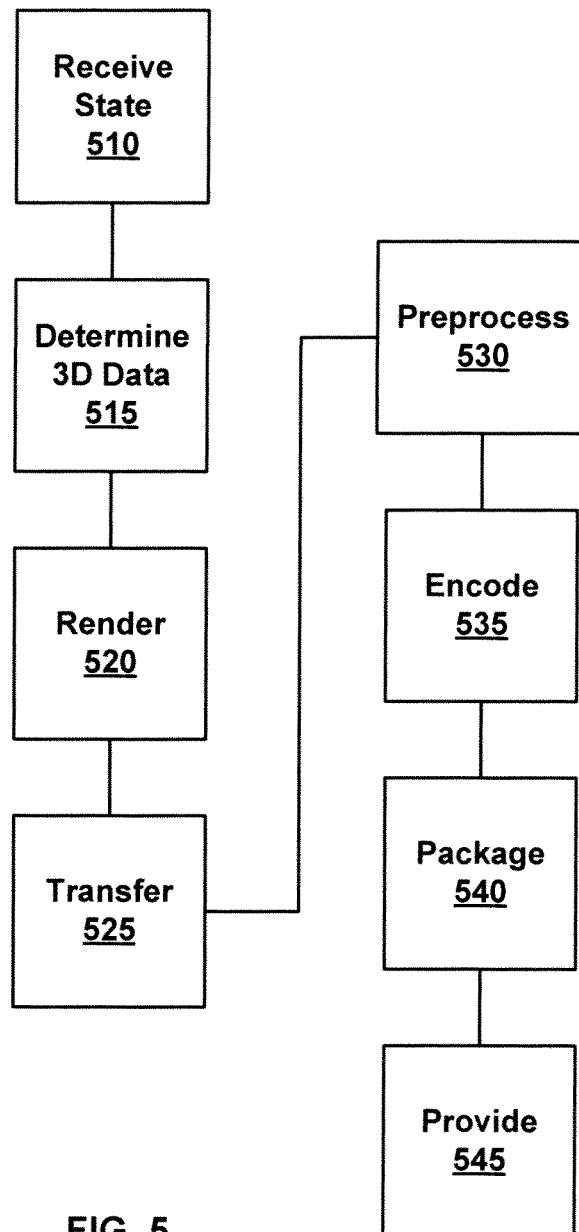


FIG. 2





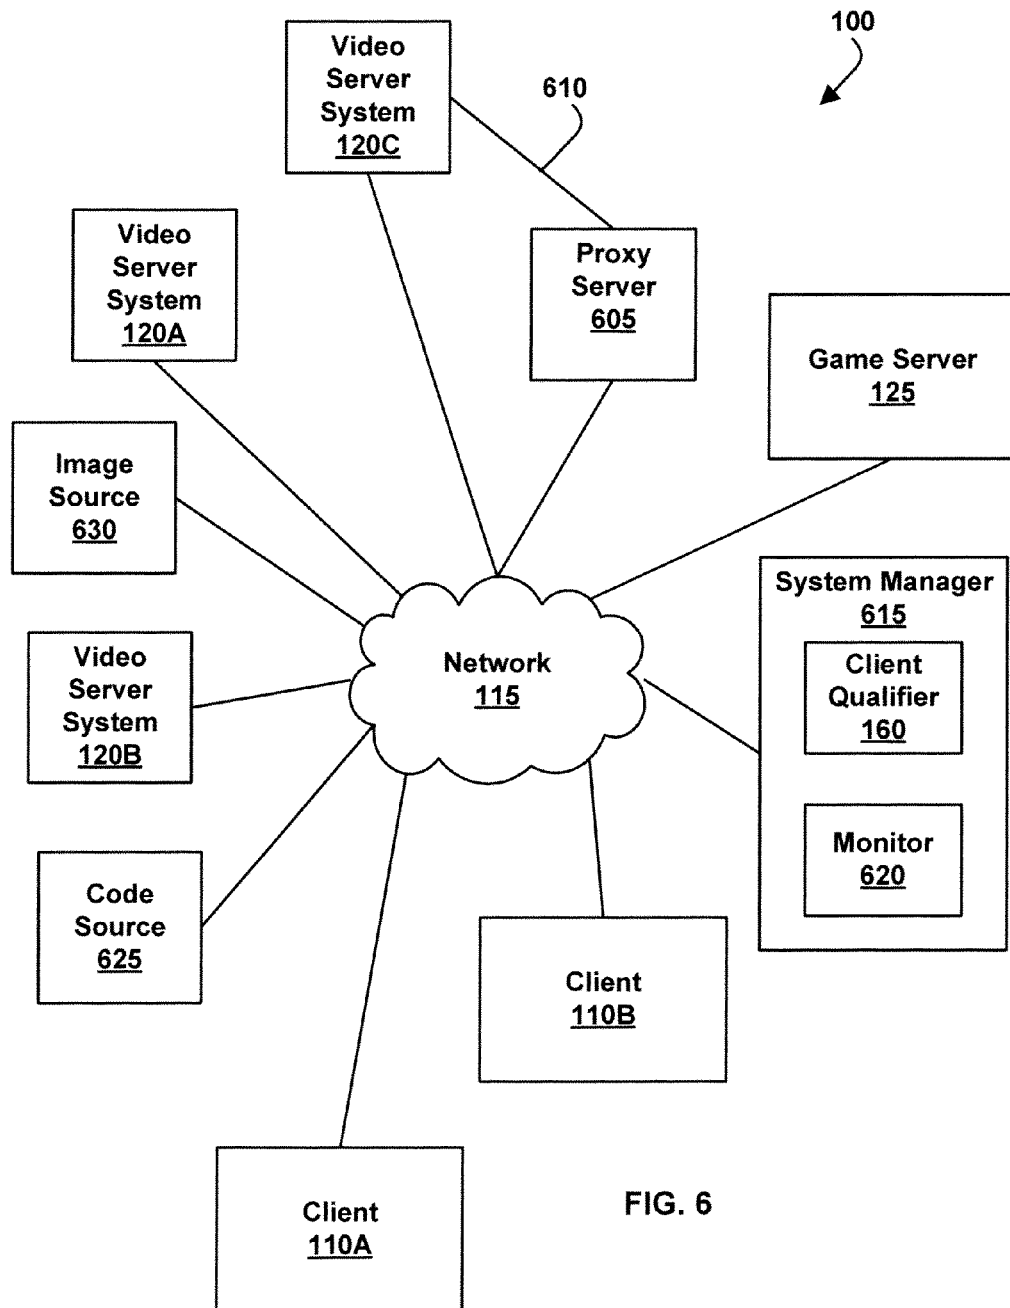


FIG. 6

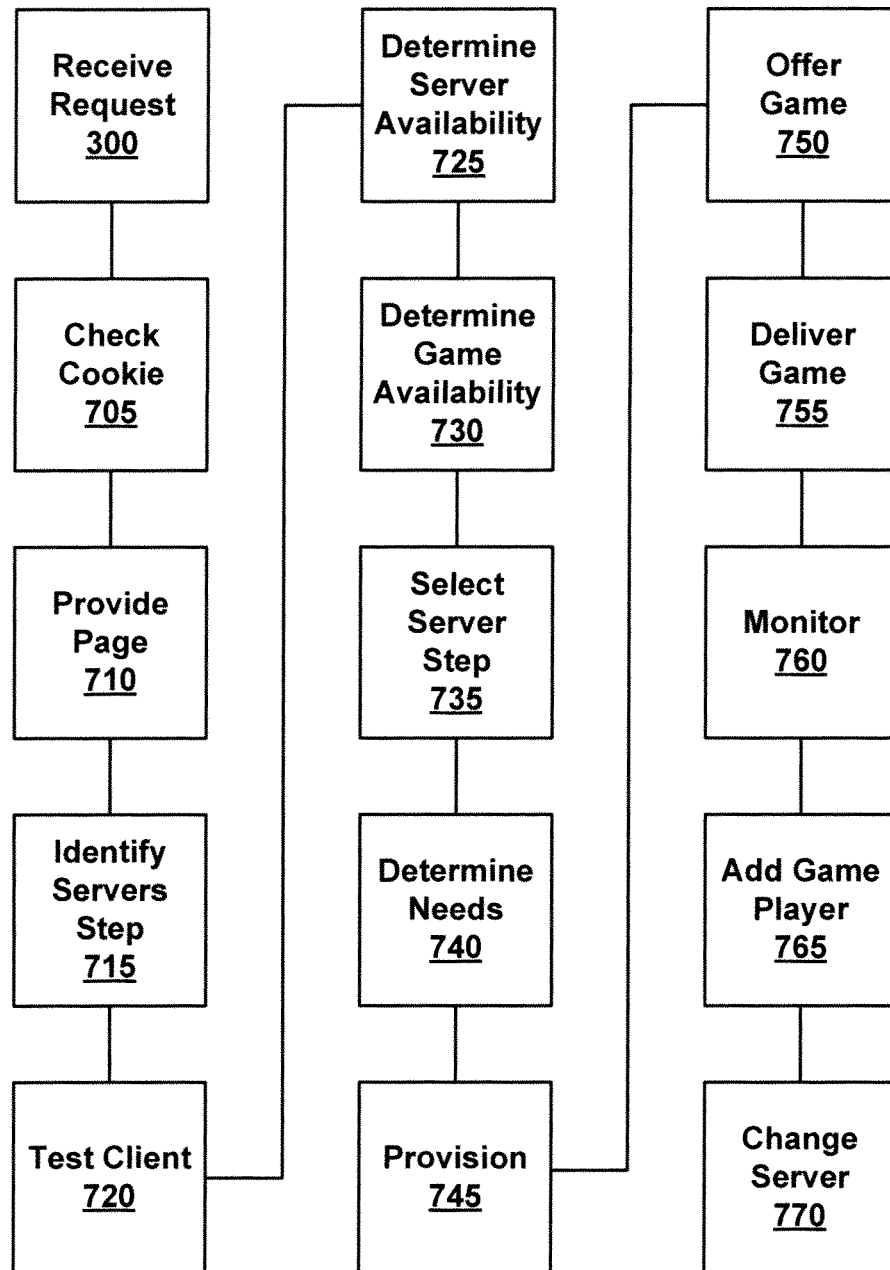


FIG. 7




FIG. 8




FIG. 9A

Chatpad maps to a regular keyboard
(red circle is mapped to Ctrl)




FIG. 9B

**Chatpad maps to a regular keyboard
(red circle is mapped to Ctrl)**




FIG. 9C




FIG. 10

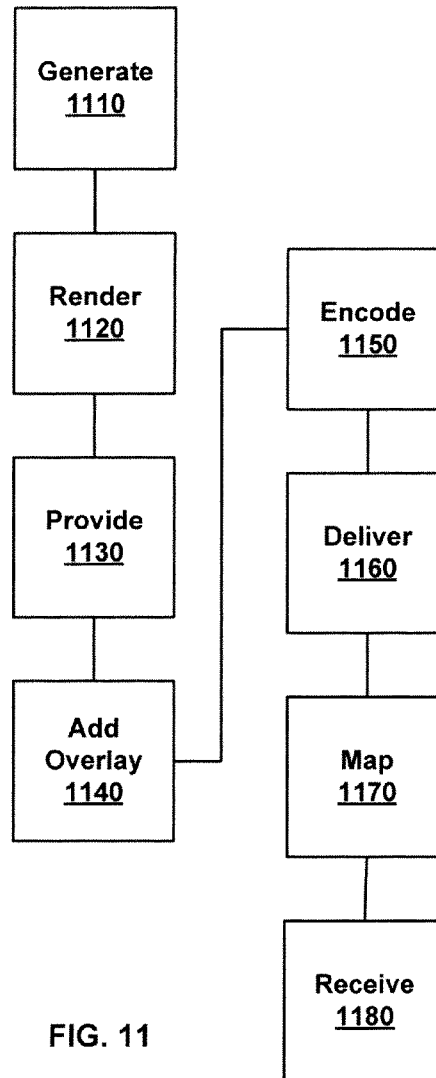


FIG. 11

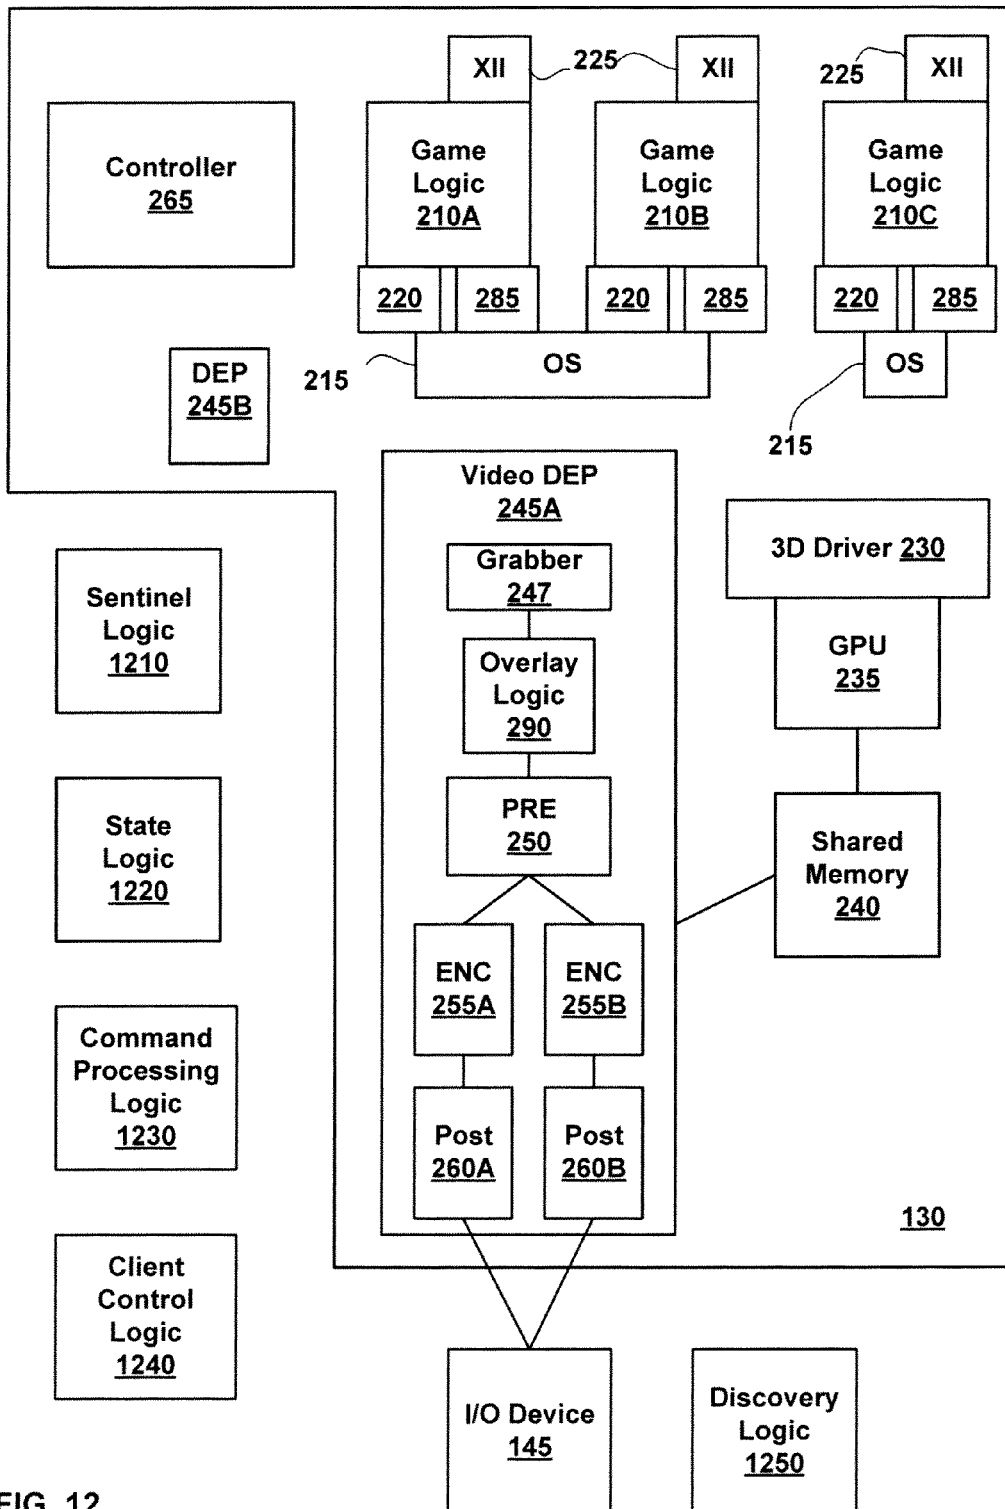




FIG. 12

**FIG. 14**

COMMAND SENTINEL

This application is a continuation-in-part of U.S. patent application Ser. No. 12/826,489 filed Jun. 29, 2010 now U.S. Pat. No. 8,888,592 and entitled "Voice Overlay," which in turn claims benefit of U.S. Provisional Patent Application Ser. No. 61/354,699 filed Jun. 14, 2010 and entitled "Video Game Overlay," and is a continuation-in-part of U.S. patent application Ser. No. 12/791,819 filed Jun. 1, 2010 and entitled "Qualified Video Delivery," which in turn claims priority and benefit of the following provisional patent applications:

61/183,035 filed Jun. 1, 2009 and entitled "Game Server Architecture,"

61/183,037 filed Jun. 1, 2009 and entitled "Bufferless H.264 Variant,"

61/183,088 filed Jun. 2, 2009 and entitled "I/O Level Virtualization,"

61/183,546 filed Jun. 2, 2009 and entitled "Self-Spawning Game Environments,"

61/323,354 filed Apr. 12, 2010 and entitled "Artificial Frames," and

61/345,534 filed May 17, 2010 and entitled "Dynamic Game Server Including Qualifier,"

further this application is a continuation-in part of U.S. patent application Ser. No. 11/499,524 filed Aug. 3, 2006 now abandoned entitled "Online Game Customer Lead Generation," and claims benefit of U.S. provisional patent applications:

61/488,297 filed May 20, 2011 and entitled "Web based Game Controller,"

61/558,376 filed Nov. 10, 2011 and entitled "Command Sentinel," and

61/558,859 filed Nov. 11, 2011 and entitled "Real-time Watermarking."

This application is related to U.S. patent applications:

Ser. No. 12/790,948 filed May 31, 2010 and entitled "Bufferless H.264 Variant," and

Ser. No. 12/790,955 filed May 31, 2010 and entitled "Game Execution Environments."

All of the above provisional and non-provisional patent applications are hereby incorporated herein by reference.

BACKGROUND**1. Field of the Invention**

The invention is in the field of audio/video processing and more specifically in the field of audio/video overlays.

2. Related Art

There are several models for the execution and rendering of video games. In the simplest model a game is executed and rendered on a computing device local to a player of the game. In another model a game state is maintained at a central server and communicated over a computing network to a client where rendering of video takes place. This model is commonly used by multi-player games in which the game state is dependent on inputs received over the computing network from multiple players of the game. In a third model a game state is maintained at a central server that also renders video for delivery to clients over a computing network as a video stream.

SUMMARY

The invention includes a video server system configured to render and provide game video to one or more clients over a computing network, such as the internet. This server is optionally configured to serve multiple video games to multiple clients. The rendered video is based on the state of one or

more games that can be maintained at the video game server system and/or received from a separate game server.

In some of the embodiments, a single graphics processing unit (GPU) is used to render multiple video streams representative of different video games, and or different viewpoints within a game. For example a single GPU may be used to render video streams for two different game players in parallel. Each of these video streams is based on different (independent) viewpoints of the game players and is typically provided to different clients used by each of the game players.

In some embodiments, a shared memory is used by a 3D video driver to store output of the GPU. This shared memory is accessed by one or more dynamic encoding pipeline. Transfer of video data to and/or from this shared memory is controlled by virtual 3D video drivers. This control allows delivery of the output of a shared GPU to a plurality of shared encoding pipelines in a coordinated manner.

Some embodiments include one or more encoding pipelines configured to be dynamically configured to produce various types of video streams. These pipelines are provisioned to meet specific needs of a client. They are optionally configured to provide more than one type of encoding and to change the encoding applied to a video stream automatically in response to a change in the needs of the client.

Various embodiments of the invention include a video server system comprising a client qualifier configured to remotely determine if a client is qualified to receive a video stream configured for display on the client, the qualification being based on abilities of the client to receive and display the video stream at at least a minimum level of quality; a video source configured to provide the video stream, and to provide the video stream only if the client is qualified to receive the video stream as determined by the client qualifier; an I/O device configured for communicating with the client via a communication network; a storage configured to store video data for use by the video source or the client qualifier; and an electronic processor configured to execute software components of the video source or the client qualifier.

Various embodiments of the invention include a game serving system comprising a first game application stored on a computer readable medium and configured to execute a first video game responsive to game commands received from a first client; a second game application stored on a computer readable medium and configured to execute a second game instance responsive to game commands received from a second client; logic configured to execute the first game instance and the second game instance in parallel; a first virtual 3D video driver configured to received a first game environment from the first game application; a second virtual 3D video driver configured to received second game environment from the second game application; a 3D video driver configured to receive an output of the first and second virtual 3D video drivers; a graphics processing unit configured to generate a first video stream by rendering the first game environment, and to generate a second video stream by rendering the second game environment, in parallel; a first encoding pipeline configured to encode the first video stream in a first codec; a second encoding pipeline configured to encode the second video stream in a second codec; and an output configured to provide the first video stream to the first client and the second video stream to the second client.

Various embodiments of the invention include a game serving system comprising a first game application stored on a computer readable medium and configured to execute a first game instance responsive to game commands received from a first client; a second game application stored on a computer readable medium and configured to execute a second game

instance responsive to game commands received from a second client; logic configured to execute the first game instance and the second game instance in parallel; a 3D video driver configured to received different game environments from the first game application and from the second game application; a graphics processing unit configured to generate a first video frame by rendering the game environment from the first game application, and to generate a second video frame by rendering the game environment from the second game application; a first encoding pipeline configured to encode the first video frame in a first video format; a second encoding pipeline configured to encode the second video frame in a second video format; a memory shared by at least the 3D video driver and at least the first encoding pipeline; logic configured to transfer the first video frame from the 3D video driver to the shared memory in sync with the production of the second video frame; and an output configured to provide the first video stream to the first client and the second video stream to the second client.

Various embodiments of the invention include a game serving system comprising a first game application stored on a computer readable medium and configured to execute a first game instance responsive to game commands received from a first client; a second game application stored on a computer readable medium and configured to execute a second game instance responsive to game commands received from a second client; logic configured to execute the first game instance and the second game instance in parallel; logic configured to provision a first encoder pipeline and a second encoder pipeline, the first encoder pipeline including a dynamically allocated encoder selected based on needs of the first client, the second encoder pipeline including a dynamically allocated encoder selected based on needs of the second client; one or more processor configured to execute the logic configured to spawn a first encoder pipeline and a second encoder pipeline, the first game application, the second game application; and an output configured to provide a first video stream generated by the first encoder pipeline to the first client, and a second video stream generated by the second encoder pipeline to the second client.

Various embodiments of the invention include a game delivery system comprising a plurality of video server systems each configured to provide streaming video to a plurality of clients over a communication network, the video server systems being geographically distributed, contents of the streaming video being dependant on game commands received from one or more of the clients and game logic included within each of the video server system; and a client qualifier configured to determine if a first of the plurality of clients is qualified to receive the streaming game video and to determine which of the video server system is best qualified to provide the streaming game video to the first of the plurality of clients. Optionally, each of the video server systems is configured to provision a dynamically encoded encoding pipeline responsive to needs of the first of the plurality of clients, and optionally the streaming video is configured to be displayed within a browser. In some embodiments, the determination of which of the video server systems is best qualified is dependent on characteristics of a communication channel between the determined video server system and the first of the plurality of clients.

Various embodiments of the invention include a method comprising receiving a request from a client to display an object, the object being configured to allow a user of the client to play a video game; testing the client to determine if the client is qualified to receive a streaming video of the game; qualifying the client based on the testing; selecting a video

server system from among a plurality of geographically distributed video server systems to provide streaming video to the client, contents of the streaming video being dependant on game commands received from the client and game logic of the video game included the selected video server system, the selection being based on which of the plurality of video server systems is best qualified to provide the streaming video to the client. Optionally, the plurality of geographically distributed video systems includes video server systems disposed in different cities. In some embodiments the method further comprises adding a game player to the game and repeating the selecting the video server system, the selection being further based on which of the plurality of video server systems is best qualified to provide streaming video to a client of the added game player, and optionally changing which of the video server systems provides streaming video to the client based on the repeated selection.

Various embodiments of the invention include a method of generating game video, the method comprising generating a game environment based on game rules and game commands received from one or more players; rendering the game environment to a video frame or audio data; providing the video frame or audio data to an encoding pipeline; adding an overlay to the video frame or audio data; encoding the overlay and the video frame or audio data; and delivering the encoded overlay and video frame or audio data to a client of the one or more players via the internet.

Various embodiments of the invention include a game delivery system comprising first game logic configured to provide audio data based on events within a video game; a first encoding pipeline configured to place an audio overlay on the audio data and to encode the audio overlay and the audio data into an audio packet according to an audio codec; an output configured to provide the audio packet to a first client over a communication network; and an electronic processor configured to execute the first game logic or the encoding pipeline.

Various embodiments of the invention include a game delivery system comprising game logic configured to provide a game environment based on game rules and game commands received from one or more game players; a graphics processing unit configured to render the game environment to video data; an encoding pipeline configured to place an overlay on the video data and to encode the overlay and the video data as a video frame; and an output configured to provide the video frame to a client over a communication network.

Various embodiments of the invention include a method of generating game video, the method comprising generating a game environment based on game rules and game commands received from one or more players; rendering the game environment to a video frame or audio data; providing the video frame or audio data to an encoding pipeline; adding an overlay to the video frame or audio data; encoding the overlay and the video frame or audio data; and delivering the encoded overlay, and video frame or audio data to a client of the one or more players via the internet.

Various embodiments of the invention include a system comprising first game logic configured to provide first audio data based on events within a video game; an audio input configured to receive second audio data from a first client; a first encoding pipeline configured to place the second audio data as an audio overlay on the first audio data, and to encode the first audio data and the second audio data into a first audio packet; an output configured to provide the first audio packet to a second client over a communication network; and an electronic processor configured to execute the first game logic or the encoding pipeline.

5

Various embodiments of the invention include a method of generating game audio, the method comprising: generating game audio based on events within a game; receiving audio data from a first player of the game over a computing network; overlaying the game audio and the audio data; encoding the game audio and the audio data to generate an audio stream; and delivering the audio stream including both the game audio and the audio data to a second player of the game over the computing network.

Various embodiments of the invention include a video server system comprising an application engine configured to execute a first computer application a graphics processing unit configured to generate a graphic output in response to graphics data received from the first computer application, an encoding pipeline configured to generate an encoded video stream using the graphic output, sentinel logic configured to monitor the graphic output and determine that a blocked input control is included within the graphic output, the blocked input control being configured to send a blocked command to the first computer application, overlay logic configured to place an overlay on the graphic output, the overlay configured to change the appearance of the blocked input control, an output configured to communicate the video stream to a remote client via a communication network, and a processor configured to execute the sentinel logic, the overlay logic or the command processing logic.

Various embodiments of the invention include a method of generating game video, the method comprising generating a game environment based on game rules of a video game and game commands received from one or more players, rendering the game environment to a video frame, providing the video frame to an encoding pipeline, monitoring the game environment or the video frame for the presence of a blocked input control, encoding the video frame to generate part of a video stream, adding an overlay to the video frame or the video stream, the overlay being configured to modify the blocked input control, and delivering the encoded video stream, including the overlay, to a client of the one or more players via the internet.

Various embodiments of the invention include a computer readable medium having stored thereon computing instructions configured to monitor a graphics output of a graphics processing unit, computing instructions configured to detect a blocked input control within the graphics output, and computing instructions configured to cause an overlay to be placed on the blocked input control so as to change an appearance of the blocked input control.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a game system, according to various embodiments of the invention.

FIG. 2 is a block diagram of a video source configured to serve multiple video games, according to various embodiments of the invention.

FIG. 3 illustrates methods of receiving a qualified request for a video stream from a client, according to various embodiments of the invention.

FIG. 4 illustrates methods of allocating resources within a video source, according to various embodiments of the invention.

FIG. 5 illustrates methods of serving multiple video games, according to various embodiments of the invention.

FIG. 6 illustrates a game system including more than one video server system, according to various embodiments of the invention.

6

FIG. 7 illustrates methods of providing a video stream to a qualified client, according to various embodiments of the invention.

FIG. 8 illustrates overlays on game video, according to various embodiments of the invention.

FIGS. 9A, 9B and 9C illustrate game control mappings, according to various embodiments of the invention.

FIG. 10 illustrates part of an input map configured for mapping of user input controls to game commands.

FIG. 11 illustrates methods of using an overlay, according to various embodiments of the invention.

FIG. 12 illustrates parts of a video server system including sentinel logic, according to various embodiments of the invention.

FIG. 13 illustrates methods of generating game video, according to various embodiments of the invention.

FIG. 14 illustrates a method of identifying blocked input controls, according to various embodiments of the invention.

DETAILED DESCRIPTION

FIG. 1 is a block diagram of a Game System 100, according to various embodiments of the invention. Game System 100 is configured to provide a video stream to one or more Clients 110 via a Network 115. Game System 100 typically includes a Video Server System 120 and an optional Game Server 125. Video Server System 120 is configured to provide the video stream to the one or more Clients 110 with a minimal quality of service. For example, Video Server System 120 may receive a game command that changes the state of or a point of view within a video game, and provide Clients 110 with an updated video stream reflecting this change in state with minimal lag time. The Video Server System 120 may be configured to provide the video stream in a wide variety of alternative video formats, including formats yet to be defined. Further, the video stream may include video frames configured for presentation to a user at a wide variety of frame rates. Typical frame rates are 30 frames per second, 60 frames per second, and 120 frames per second. Although higher or lower frame rates are included in alternative embodiments of the invention.

Clients 110, referred to herein individually as 110A, 110B, etc., may include terminals, personal computers, game consoles, tablet computers, telephones, set top boxes, kiosks, wireless devices, digital pads, stand-alone devices, handheld game playing devices, and/or the like. Typically, Clients 110 are configured to receive encoded video streams, decode the video streams, and present the resulting video to a user, e.g., a player of a game. The processes of receiving encoded video streams and/or decoding the video streams typically includes storing individual video frames in a receive buffer of the client. The video streams may be presented to the user on a display integral to Client 110 or on a separate device such as a monitor or television. Clients 110 are optionally configured to support more than one game player. For example, a game console may be configured to support two, three, four or more simultaneous players. Each of these players may receive a separate video stream, or a single video stream may include regions of a frame generated specifically for each player, e.g., generated based on each player's point of view. Clients 110 are optionally geographically dispersed. The number of clients included in Game System 100 may vary widely from one or two to thousands, tens of thousands, or more. As used herein, the term "game player" is used to refer to a person that plays a game and the term "game playing device" is used to refer to a device used to play a game.

Clients **110** are configured to receive video streams via Network **115**. Network **115** may be any type of communication network including, a telephone network, the Internet, wireless networks, powerline networks, local area networks, wide area networks, private networks, and/or the like. In typical embodiments, the video streams are communicated via standard protocols, such as TCP/IP or UDP/IP. Alternatively, the video streams are communicated via proprietary standards.

A typical example of Clients **110** is a personal computer comprising a processor, non-volatile memory, a display, decoding logic, network communication capabilities, and input devices. The decoding logic may include hardware, firmware, and/or software stored on a computer readable medium. Systems for decoding (and encoding) video streams are well known in the art and vary depending on the particular encoding scheme used.

Clients **110** may, but are not required to, further include systems configured for modifying received video. For example, a client may be configured to perform further rendering, to overlay one video image on another video image, to crop a video image, and/or the like. For example, Clients **110** may be configured to receive various types of video frames, such as I-frames, P-frames and B-frames, and to process these frames into images for display to a user. In some embodiments, a member of Clients **110** is configured to perform further rendering, shading, conversion to 3-D, or like operations on the video stream. A member of Clients **110** is optionally configured to receive more than one audio or video stream. Input devices of Clients **110** may include, for example, a keyboard, a joystick, a pointing device, a force feedback device, a motion and/or location sensing device, a mouse, a touch screen, a neural interface, a camera, input devices yet to be developed, and/or the like.

The video stream (and optionally audio stream) received by Clients **110** is generated and provided by Video Server System **120**. As is described further elsewhere herein, this video stream includes video frames (and the audio stream includes audio frames). The video frames are configured (e.g., they include pixel information in an appropriate data structure) to contribute meaningfully to the images displayed to the user. A meaningful contribution is a contribution that is readily observable by a game player. As used herein, the term “video frames” is used to refer to frames including predominantly information that is configured to contribute to, e.g. to effect, the images shown to the user. Most of the teachings herein with regard to “video frames” can also be applied to “audio frames.”

Clients **110** are typically configured to receive inputs from a user. These inputs may include game commands configured to change the state of the video game or otherwise affect game play and/or operation of a video game. The game commands can be received using input devices and/or may be automatically generated by computing instructions executing on Clients **110**. The received game commands are communicated from Clients **110** via Network **115** to Video Server System **120** and/or Game Server **125**. For example, in some embodiments, the game commands are communicated to Game Server **125** via Video Server System **120**. In some embodiments, separate copies of the game commands are communicated from Clients **110** to Game Server **125** and Video Server System **120**. The communication of game commands is optionally dependent on the identity of the command. Game commands are optionally communicated from Client **110A** through a different route or communication channel than that used to provide audio or video streams to Client **110A**.

Game Server **125** is optionally operated by a different entity than Video Server System **120**. For example, Game Server **125** may be operated by the publisher of a multiplayer game. In this example, Video Server System **120** is optionally viewed as a client by Game Server **125** and optionally configured to appear from the point of view of Game Server **125** to be a prior art client executing a prior art game engine. Communication between Video Server System **120** and Game Server **125** optionally occurs via Network **115**. As such, Game Server **125** can be a prior art multiplayer game server that sends game state information to multiple clients, one of which is Video Server System **120**. Video Server System **120** may be configured to communicate with multiple instances of Game Server **125** at the same time. For example, Video Server System **120** can be configured to provide a plurality of different video games to different users. Each of these different video games may be supported by a different Game Server **125** and/or published by different entities. In some embodiments, several geographically distributed instances of Video Server System **120** are configured to provide game video to a plurality of different users. Each of these instances of Video Server System **120** may be in communication with the same instance of Game Server **125**. Communication between Video Server System **120** and one or more Game Server **125** optionally occurs via a dedicated communication channel. For example, Video Server System **120** may be connected to Game Server **125** via a high bandwidth channel that is dedicated to communication between these two systems.

Video Server System **120** comprises at least a Video Source **130**, an I/O Device **145**, a Processor **150**, and non-transitory Storage **155**. Video Server System **120** may include one computing device or be distributed among a plurality of computing devices. These computing devices are optionally connected via a communications system such as a local area network.

Video Source **130** is typically configured to provide a video stream, e.g., streaming video or a series of video frames that form a moving picture. In some embodiments, Video Source **130** is also configured to provide an audio stream. In some embodiments, Video Source **130** includes a video game engine and rendering logic. The video game engine is configured to receive game commands from a player and to maintain a copy of the state of the video game based on the received commands. This game state includes the position of objects in a game environment, as well as typically a point of view. The game state may also include properties, images, colors and/or textures of objects. The game state is typically maintained based on game rules, as well as game commands such as move, turn, attack, set focus to, interact, use, and/or the like. Part of the game engine is optionally disposed within Game Server **125**. Game Server **125** may maintain a copy of the state of the game based on game commands received from multiple players using geographically disperse clients. In these cases, the game state is provided by Game Server **125** to Video Source **130**, wherein a copy of the game state is stored and rendering is performed. Game Server **125** may receive game commands directly from Clients **110** via Network **115**, and/or may receive game commands via Video Server System **120**.

In alternative embodiments, Video Source **130** is configured to provide images that do not form a moving picture. These images may be the image output of a game application or, as discussed elsewhere herein, the image output of a non-game application.

Video Source **130** typically includes rendering logic, e.g., hardware, firmware, and/or software stored on a computer

readable medium such as Storage **155**. This rendering logic is configured to create video frames of the video stream based on the game state. All or part of the rendering logic is optionally disposed within a graphics processing unit (GPU). Rendering logic typically includes processing stages configured for determining the three-dimensional spatial relationships between objects and/or for applying appropriate textures, etc., based on the game state and viewpoint. The rendering logic produces raw video that is then usually encoded prior to communication to Clients **110**. For example, the raw video may be encoded according to an Adobe Flash® standard, .wav, H.264, H.263, On2, VP6, VC-1, WMA, Huffvuv, Lagarith, MPG-x, Xvid, FFmpeg, x264, VP6-8, realvideo, mp3, or the like. The encoding process produces a video stream that is optionally packaged for delivery to a decoder on a remote device. The video stream is characterized by a frame size and a frame rate. Typical frame sizes include 800×600, 1280×720 (e.g., 720p), 1024×768, although any other frame sizes may be used. The frame rate is the number of video frames per second. A video stream may include different types of video frames. For example, the H.264 standard includes a “P” frame and an “I” frame. I-frames include information to refresh all macro blocks/pixels on a display device, while P-frames include information to refresh a subset thereof. P-frames are typically smaller in data size than are I-frames. As used herein the term “frame size” is meant to refer to a number of pixels within a frame. The term “frame data size” is used to refer to a number of bytes required to store the frame. Rendering logic is optional in embodiments where the output of Video Source **130** does not include images generated in from a game environment or the like.

In alternative embodiments Video Source **130** includes a video recording device such as a camera. This camera may be used to generate delayed or live video that can be included in the video stream of a computer game. The resulting video stream optionally includes both images rendered in real-time and images recorded using a still or video camera. Video Source **130** may also include storage devices configured to store previously recorded video to be included in a video stream. Video Source **130** may also include motion or positioning sensing devices configured to detect motion or position of an object, e.g., person, and logic configured to determine a game state or produce video-based on the detected motion and/or position.

Video Source **130** is optionally configured to provide overlays configured to be placed on other video. For example, these overlays may include a command interface, log in instructions, messages to a game player, images of other game players, video feeds of other game players (e.g., webcam video). In embodiments of Client **110A** including a touch screen interface, the overlay may include a virtual keyboard, joystick, touch pad, and/or the like. In one example of an overlay a player’s voice is overlaid on an audio stream. Video Source **130** optionally further includes one or more audio sources.

As discussed elsewhere herein, overlays are optionally used to cover or otherwise modify elements within a video image generated by a game engine. For example, an input control that is undesirable when a video game is executed on a Video Server System **120** (but perhaps not undesirable when executed locally on a personal computer) can be covered using an overlay. An input control is a clickable, selectable or otherwise actionable area of an image provided by Video Server System **120** that is tied to underlying functionality. The underlying functionality optionally includes generation of an input command, such as a game command. Not all input commands are generated by an input control. For example,

some input commands are generated using keystrokes. Input commands may be configured to operate on the local member of Clients **110** or be configured to operate on software or hardware that happens to be remote from the local member of Clients **110**. For example, an input command may be configured to operate on a game engine within Video Source **130**. In some embodiments, input commands include game commands as well as other commands initiated by a game player.

In embodiments wherein Video Server System **120** is configured to maintain the game state based on input from more than one player, each player may have a different point of view comprising a position and direction of view. Video Source **130** is optionally configured to provide a separate video stream for each player based on their point of view. Further, Video Source **130** may be configured to provide a different frame size, frame data size, and/or encoding to each of Client **110**. Video Source **130** is optionally configured to provide 3-D video.

I/O Device **145** is configured for Video Server System **120** to send and/or receive information such as video, commands, requests for information, a game state, client identities, player identities, game commands, security information, audio data, and/or the like. I/O Device **145** typically includes communication hardware such as a network card or modem. I/O Device **145** is configured to communicate with Game Server **125**, Network **115**, and/or Clients **110**. I/O Device **145** is configured to receive the information from more than one of Clients **110**. I/O Device **145** is optionally configured to receive the information as packets using a standard such as TCP or UDP.

Processor **150** is configured to execute logic, e.g. software, included within the various components of Video Server System **120** discussed herein. For example, Processor **150** may be programmed with software instructions in order to perform the functions of Video Source **130**, Game Server **125**, and/or a Client Qualifier **160**. Video Server System **120** optionally includes more than one instance of Processor **150**. Processor **150** may also be programmed with software instructions in order to execute commands received by Video Server System **120**, or to coordinate the operation of the various elements of Game System **100** discussed herein. Processor **150** may include one or more hardware device. Processor **150** is an electronic or optical processor.

Storage **155** includes non-transitory analog and/or digital storage devices. For example, Storage **155** may include an analog storage device configured to store video frames. Storage **155** may include a computer readable digital storage, e.g. a hard drive, an optical drive, or solid state storage. Storage **115** is configured (e.g. by way of an appropriate data structure or file system) to store video frames, artificial frames, a video stream including both video frames and artificial frames, audio frame, an audio stream, overlays, command mappings, lists of blocked commands, bitmaps, and/or the like. Storage **155** is optionally distributed among a plurality of devices. In some embodiments, Storage **155** is configured to store the software components of Video Source **130** discussed elsewhere herein. These components may be stored in a format ready to be provisioned when needed.

Video Server System **120** optionally further comprises Client Qualifier **160**. Client Qualifier **160** is configured for remotely determining the capabilities of a client, such as Clients **110A** or **110B**. These capabilities can include both the capabilities of Client **110A** itself as well as the capabilities of one or more communication channels between Client **110A** and Video Server System **120**. For example, Client Qualifier **160** may be configured to test a communication channel through Network **115**.

11

Client Qualifier **160** can determine (e.g., discover) the capabilities of Client **110A** manually or automatically. Manual determination includes communicating with a user of Client **110A** and asking the user to provide capabilities. For example, in some embodiments, Client Qualifier **160** is configured to display images, text, and/or the like within a browser of Client **110A**. The displayed objects request that the user enter information such as operating system, processor, video decoder type, type of network connection, display resolution, etc. of Client **110A**. The information entered by the user is communicated back to Client Qualifier **160**.

Automatic determination may occur, for example, by execution of an agent on Client **110A** and/or by sending test video to Client **110A**. The agent may comprise computing instructions, such as java script, embedded in a web page or installed as an add-on. The agent is optionally provided by Client Qualifier **160**. In various embodiments, the agent can find out processing power of Client **110A**, decoding and display capabilities of Client **110A**, lag time reliability and bandwidth of communication channels between Client **110A** and Video Server System **120**, a display type of Client **110A**, firewalls present on Client **110A**, hardware of Client **110A**, software executing on Client **110A**, registry entries within Client **110A**, and/or the like.

Client Qualifier **160** includes hardware, firmware, and/or software stored on a computer readable medium. Client Qualifier **160** is optionally disposed on a computing device separate from one or more other elements of Video Server System **120**. For example, in some embodiments, Client Qualifier **160** is configured to determine the characteristics of communication channels between Clients **110** and more than one instance of Video Server System **120**. In these embodiments the information discovered by Client Qualifier can be used to determine which instance of Video Server System **120** is best suited for delivery of streaming video to one of Clients **110**.

FIG. 2 is a block diagram of embodiments of Video Source **130** configured to serve multiple video games. The multiple video games can include multiple instances of the same video game and/or instances of different video games. The video games are optionally multiplayer games. For example, a game state of one or more of the video games may be maintained by one or more instances of Game Server **125** each based on inputs received from multiple clients. The elements of Video Source **130** illustrated in FIG. 2 are optionally executed using Processor **150**.

The video games are executed using different instances of Game Logic **210**, individually labeled **210A**, **210B**, **210C**, etc. Game Logic **210** is configured to use a game state to determine a game environment that can be rendered to a video stream configured to be presented to a game player. The game environment is a three dimensional virtual environment including in-game objects, avatars, locations of objects, their shapes, textures, and spatial relationships there between, and the like. A game environment can include vertex data, transformation data and texture data, and/or the like. In alternative embodiments one or more of Game Logic **210** are replaced by other application logic. In these embodiments the output of Video Source **130** can include video, single images, and/or multiple images at a variable frame rate. Frame rate may vary, for example, as a function of how often an image is updated by an application. In some embodiments images are provided in a video format but new frames are typically only provided when the generated image changes.

The rendering of the game environment is typically based on one or more point of view associated with a specific game player. Video Source **130** may include more than 1, 2, 3, 5, 7,

12

15, or 31 instances of Game Logic **210**. Game Logic **210** is optionally a client of Game Server **125** and may be configured to communicate with Game Server **125** via Network **115**.

Game Logic **210** is configured to receive game commands from one or more of Clients **110** and to process the received commands according to a set of game rules. These rules cover, for example, how avatars interact with other avatars or in game objects, avatar movement, game instance management, and/or the like. Game Logic **210** is optionally also configured to generate audio data based on events within the game. This audio data may represent a gunshot, a splash, an engine, voice, flying, rain, music, or any other sound that could occur in a game. For example, an event such as one object hitting another may result in audio data representing a related sound. Game Logic **210** includes hardware, firmware, and/or software stored on a computer readable medium.

Each instance of Game Logic **210** can be disposed on a separate computing device or several instances of Game Logic **210** could be disposed on the same computing device or a single instance of Game Logic **210** could be disposed across multiple computing devices. Game instances can dynamically change the number and/or identify of computing devices used depending on the requirements of the game environment the user is currently experiencing. Instances of Game Logic **210** disposed on the same computing device are optionally executed within separate virtual machines or virtual I/O shells. In some embodiments, different instances of Game Logic **210** are configured to communicate game commands and/or game state information directly to each other, e.g., without necessarily communication through Game Server **125**.

Game Logic **210** typically executes on top of an Operating System (OS) **215**. Operating System **215** may include Windows™, Linux, UNIX, Mac OS™, Solaris™ and/or the like may operate between Operating System **215** and Game Logics **210**. Virtual machine platforms such as ESX, Hyper-V, and/or the like. In these embodiments, one or more of Game Logic **210** can be executed within a virtual machine. Multiple instances of Game Logic **210** may execute on the same instance of Operating System **215**. For example, FIG. 2 illustrates Game Logic **210A** and Game Logic **210B** both executing on the same Operating System **215**. Instances of Game Logic **210** executing on the same Operating System **215** may, or may not, be configured for playing the same video game. For example, Game Logic **210A**, **210B** and **210C** may all be World of Warcraft® clients, or may be clients of World of Warcraft®, Eve® and Call to Arms® respectively.

In some embodiments, the game environments determined by Game Logic **210** are passed to an optional Virtual 3D Video Driver **220**. Virtual 3D Video Driver **220** is configured to appear, from the point of view of Game Logic **210**, as a non-virtual 3D video driver controlling a graphics processing unit. Each instance of Game Logic **210** may be associated with its own instance of Virtual 3D Video Driver **220**, or Virtual 3D Video Driver **220** may be shared by two or more instances of Game Logic **210**. In some embodiments, game audio determined by each Game Logic **210** is passed to an optional Virtual Audio Driver **285**. Game audio includes sound data configured to be presented to a game player as part of an audio stream. The game audio can result from rules of the game, e.g., shooting a gun should result in a popping sound, or may be received from other sources such as music tracks, members of Clients **110**, a microphone, and/or the like.

In alternative embodiments, Game Logic **210** is configured to include the functionality, as discussed herein, of Virtual 3D Video Driver **220** and/or Virtual Audio Driver **285**. In these

13

embodiments, Virtual 3D Video Driver **220** and/or Virtual Audio Driver **285** are not required.

Virtual 3D Video Driver **220** is further configured to pass the received game environments to a (non-virtual) 3D Driver **230**. Optionally the delivery of game environments to 3D Driver **230** is coordinated by the various instances of Virtual 3D Video Driver **220**. For example, delivery can be coordinated such that 3D Driver **230** receives game environments from only one or a minimum number of Virtual 3D Video Driver **220** at a time. In typical embodiments, each of Virtual 3D Video Drivers **220** is configured such that they appear to be a separate process and a separate source of video data to 3D Driver **230**. As such, 3D Driver **230** is configured to keep track of which video data results in which video frames after rendering.

The video data received by 3D Driver **230** are passed to Graphics Processing Unit (GPU) **235** for rendering into raw video frames. Graphics Processing Unit **235** is optionally used to render more than one video stream in parallel. For example, Graphics Processing Unit **235** may generate a first video stream based on video data received from Game Logic **210A**, generate a second video stream based on video data received from Game Logic **210B**, generate a third video stream based on video data received from Game Logic **210C**, etc. in parallel. The parallel production of video streams includes the generation of these streams at the same time. However, parallel production of video streams may, but does not necessarily, include the processing of individual frames at the same time within Graphics Processing Unit **235**. For example, in some embodiments 3D Driver **230** alternatively passes the video data generated by the various members Game Logic **210** to Graphics Processing Unit **235**. Data generated by Game Logic **210A** is used to make a video frame, and subsequently data generated by Game Logic **210B** is used to make a video frame, etc. In this case, the video streams are produced in parallel while individual frames are produced in series. Alternatively, more than one video frame may be generated within Graphics Processing Unit **235** at the same time. In this case a first part of Graphics Processing Unit **235** is used to generate one frame while a second part of Graphics Processing Unit **235** is used to generate a second frame, the first and second frames arising from video data produced by different Game Logic **210**. In some embodiments, one set of graphical computation cores is used to generate the first frame while a second set of graphical computation cores is used to generate the second frame, at the same time. The resulting video frames are passed back to the control of 3D Driver **230**.

Virtual 3D Video Drivers **220** are typically configured to manage the transfer of raw rendered video frames from 3D Driver **230**. For example, Virtual 3D Video Drivers **220** may be used to coordinate the transfer of video frames from the control of 3D Driver **230** to a Shared Memory **240**. (Following rendering, the video frames are in a memory of Graphics Processing Unit **235** or a memory managed by 3D Driver **230**.) As Virtual 3D Video Drivers **220** manages the communication of video data and frames to and from 3D Driver **230**, in some embodiments, they are capable of placing the video frames within parts of Shared Memory **240** associated with specific Video Dynamic Encoding Pipelines (DEPs) **245**. The Video DEPs **245** are individually identified as Video DEP **245A**, Video DEP **245B**, Video DEP **245C**, etc. In these embodiments, each Video DEP **245** is assigned specific memory locations within Shared Memory **240**, and is configured to retrieve video data from those locations.

In other embodiments, Virtual 3D Video Drivers **220** are configured to manage the transfer of video data to Shared

14

Memory **240** based on timing. In these embodiments, the transfers managed by each of Virtual 3D Video Drivers **220** are synchronized and the Virtual 3D Video Drivers **220** notify each Video DEP **245** when data for them is in Shared Memory **240**. Once this data is retrieved by the notified member of Video DEPs **245**, data for another member of DEPs is transferred from 3D Driver **230** to Shared Memory **240** under the control of Virtual 3D Video Drivers **220**. In these cases pixel read back from the local GPU memory to Shared Memory **240** can be based on a schedule. As used herein the terms “in sync” and “synchronized” are meant to mean that the two events are related in time by a schedule, by a timing signal, by a time delay, or that one event does not start until the other event has satisfied a condition, e.g., has completed. For example, Virtual 3D Video Drivers **220** may operate in sync such that a game environment is provided to 3D Driver **230** from a first of Virtual 3D Video Drivers **220** once Graphics Processing Unit **235** has completed rendering of a game environment from a second of Virtual 3D Video Drivers **220**. The timing of this synchronization is optionally selected to make optimum use of Graphics Processing Unit **235**.

As a result of the management by Virtual 3D Video Drivers **220**, the multiple video streams can be stored in Shared Memory **240** without confusing which frames belong to which video stream. Shared Memory **240** is optionally configured to store audio as well as video data. This audio data may be stored and/or managed in ways similar to that discussed with respect to video data.

Virtual Audio Drivers **285** are optionally part of a virtual I/O shell between Game Logic **210** and OS **215**. They are configured to appear from the point of view of Game Logic **210** as an audio driver and to pass any audio data they receive from Game Logic **210** to one of Audio DEP **270**. For example, in some embodiments, the Virtual Audio Driver **285** associated with Game Logic **210** is configured to send audio data to Audio DEP **270A**, and the Virtual Audio Driver **285** associated with Game Logic **210B** is configured to send audio data to Audio DEP **270B**, etc.

Shared Memory **240** includes random access memory (RAM) or a similar memory configured for efficient reading and writing of video data. Shared Memory **240** is configured to store video data for a plurality of different Video DEP **245**. Video data for different Video DEPs **245** is optionally stored at the same time in Shared Memory **240**. Shared Memory **240** may consist of a single hardware device or may include multiple devices.

Video DEPs **245** are dynamically allocated encoding pipelines that are each configured to encode video data rendered using Graphics Processing Unit **235**. Each member of Video DEPs **245** is configured to encode to video formats specified at the time the Video DEP **245** is provisioned. This format specification is typically based on the needs of one of Clients **110** and/or the capabilities of the communication path between Video Server System **120** and the Client **110**. Video DEPs **245** are optionally provisioned dynamically in response from a request from one of Clients **110**. For example, when Client **110A** connects to Video Server System **120** and sends a request for a video stream, Video DEP **245** can be provisioned to include elements, such as an encoder, selected specifically to meet needs of Client **110A**. As is discussed elsewhere herein, a member of Video DEPs **245** is optionally configured to use more than one encoding scheme.

Video DEPs **245** each include a Grabber **247** configured to grab video data from Shared Memory **240** and transfer this video data to a memory of the Video DEP **245**. The Grabber **247** is optionally under the control of a member of Virtual 3D Video Drivers **220**. For example, in various embodiments,

15

Grabber **247** is controlled by a member of Virtual 3D Video Drivers **220** to grab the video data from a specific location within Shared Memory **240** and/or at a specific time in synchronization with the transfer of video data to other members of Video DEPs **245**. In some embodiments, Grabber **247** is also configured to grab 2D video data directly from Memory XII **225**.

Video DEP **245** each optionally include an Overlay Logic **290** configured to overlay one or more video images on the video data, e.g., frames, generated by Graphics Processing Unit **235**. An overlay is a first image to be placed on a second image, or a sound added to another sound. Overlays may be applied with varying levels of transparency. For example, the first image can be opaque and, thus, completely hide covered parts of the second image, or the first image can be partially transparent in which case the second image can still be somewhat visible under the first image. An overlay may cover all or part(s) of the underlying image. For example, an overlay may be configured to cover certain pixels within a larger area. In this case the overlay is mapped to pixels on a display of one of Clients **110** or mapped to pixels within an image generated by Graphics Processing Unit **235**. The mapping can be accomplished in various ways. For example, a lower left pixel of the overlay may be assigned to a specific pixel on a video frame generated by Graphics Processing Unit **235**. In some embodiments, Overlay Logic **290** is configured to resize an overlay prior to applying the overlay. Overlay Logic **290** is optionally located in other positions within Video DEP **245A**. For example, after Preprocessor **250**. Overlay Logic **290** is optionally configured to apply more than one overlay. In some embodiments, Video DEP **245** includes a plurality of Overlay Logic **290** configured to apply overlays to the same or different frames. A second overlay can be placed on top of a first overlay.

The overlays can include a wide variety of pixel/video images. For example, an overlay can include a real-time or static image of a game player received via the internet, a real-time or static image of a sporting (or other) event, an image of an input control (e.g., a button, a link, a virtual joystick, a slide bar, a wheel, a clickable object, "a," "b," "TAB," "Return," "Space Bar," "Function Key," "Arrow," and/or other keys or input devices), a map, text, and/or the like. Overlays can vary widely in size and shape. In some cases the overlay is the result of rendering a 3D game environment. More than one overlay can be placed on an image. These overlays can be overlapping or separate. Overlays are optionally received in an encoded format and decoded prior to being applied. Overlays are optionally stored in Storage **155**.

The presence and content of overlays are optionally responsive to received game commands, identity of one or more game players, an account type, identity of a game being played, video captured from a prior or real-time sporting event or real-world game, game rules, and/or the content of video generated by Game Logic **210**. For example, an overlay may include a menu resulting from receipt of a game command. An overlay may include content responsive to a pixel pattern, e.g., image, included within video generated by Game Logic **210**. Overlay Logic **290** is optionally configured to apply multiple overlays to a single video frame and these overlays may be from different sources.

Overlay Logic **290** is optionally configured to apply an overlay after a video frame has been preprocessed using Preprocessor **250** and/or after encoding using one of Encoders **255**. In some embodiments the overlay includes an image received at an input of Video Source **130** (e.g., I/O Device **145**) via the internet and Overlay Logic **290** includes video processing logic configured to generate the overlay from the

16

image. The received image is optionally received from an image source discussed elsewhere herein, e.g., and Image Source **630** discussed in reference to FIG. 6.

Video DEPs **245** each optionally also include a Preprocessor (PRE) **250**. Preprocessor **250** is configured to perform a color space conversion such as RGB to YUV and/or a scaling operation to increase or decrease the resolution of the video frame. Preprocessor **250** is optional in embodiments wherein the output of Graphics Processing Unit **235** is in the YUV color space or some other desired color space. Multiple Preprocessor **250** may be included in a Video DEP **245** configured to produce multiple video streams having video frames of different sizes.

Video DEPs **245** each include at least one Encoder (ENC) **255**. Encoders **255** are individually identified as Encoder **255A**, Encoder **255B**, etc. Each of Encoders **255** is configured to encode the video data according to a specific codec, and optionally a specific color depth and/or frame size. For example, Encoders **255** may be configured to encode video data (or still images) to, a custom format, Adobe Flash® standard, .flv, .wav, .avi, .mpg, H.264, H.263, On2, VP6, VC-1, TIFF, JPEG, WMF, GIF, and/or other codecs discussed herein.

A member of Video DEPs **245** may include one, two or more Encoders **255**. These encoders may be configured to encode to different codecs and/or the different formats of the same codec. For example, Encoder **255A** may be configured to encode according to a Flash standard at a first frame size and color depth while Encoder **255** is configured to encode to the same Flash standard at a second frame size and color depth. The identity of Encoders **255** within each member of Video DEPs **245** is typically determined at the time the Video DEP **245** is provisioned. For example, a single command or command set may be used to create (provision) Video DEP **245A** and specify which components the created Video DEP **245A** should include. The creation of Video DEP **245A** is discussed further elsewhere herein. A member of Video DEPs **245** including two or more Encoders **255** is alternatively viewed as two or more separate Video DEPs **245** having some but not all components in common. For example, Video DEP **245A** and Video DEP **245B** may have the same Preprocessor **250** but different Encoders **255**.

In one example, Encoder **255A** is configured to use a codec for H.264 while Encoder **255B** is configured to use a codec for H.263. Having two or more different encoders available enables Video DEP **245A** to change encoding during delivery of a video stream. The change in encoding can be from one type of encoding to another, or merely a change in characteristics of a specific type of coding. For example, the characteristics may change in terms of color depth, number of frames per second, encoding options, number of pixels, and/or the like. In some embodiments, Video DEP **245A** is optionally configured to switch between Encoder **255A** and **255B** in response to a change in the characteristics of Client **110A** or the communication channel between Client **110A** and Video Source **130**.

In practice, when a different codec other than that being used is required a new Video DEP **245** is spawned and executed in parallel with the current Video DEP **245** for a short period of time. The new Video DEP **245** is optionally a branch of the original Video DEP **245**. For example, some components of the original Video DEP **245** may be used in the new Video DEP **245**. These two Video DEP **245** may have components logically arranged in a tree structure.

In some embodiments, Video DEP **245A** is configured to use two or more different encoders to generate two or more different video streams at the same time. These video streams

17

are based on the same game environment rendered by Graphics Processing Unit **235** and, thus, include essentially the same materials (with the possible exception of overlays) and can be sent to different places. For example, one of the video streams can be sent to Client **110A** while the other is sent to Client **110B**. Alternatively, one of the video streams can be sent to Client **110A** and the other to a website where third parties can watch the video. This website is optionally part of a social networking site or a game player's site. The two different video streams may be different in the frame rate, encoding type, frame size, color depth, etc. For example a video stream delivered to a social networking website can be of much lower quality than the video stream delivered to a game player who is playing a game using Client **110A**. The second video stream may be directed to a game player who is playing the game or to people who are merely observing the game play. A video stream is optionally directed to more than one place.

Video DEP **245A** optionally includes one or more Post Processors (Post) **260**. Individual examples of Post Processors **260** are labeled **260A** and **260B**. Post Processors **260** are configured to package an encoded video frame in a container in a format appropriate for communication over a network according to a public or proprietary network protocol. For example, some protocols such as Adobe RTMP require post processing while other video standards such as H.264 Annex B do not require post processing. Each of Post Processors **260** may be associated with a specific member of Encoders **255**, or several Post Processors **260** may be configured to receive encoded video frames from one member of Encoders **255**.

The output of Post Processors **260** is directed to I/O Device **145** for delivery to one or more of Clients **110**. The Elements of Video DEPs **245** discussed herein include hardware, firmware and/or software stored on a computer readable medium. For example, each of Video DEPs **245** may represent a set of software loaded into memory and executing using an electronic processor.

The operation of Video Source **130** is typically managed by a Controller **265**. Controller **265** includes hardware, firmware and/or software stored on a computer readable medium. For example Controller **265** may include software stored in memory and executed using a microprocessor.

Audio DEPs **270** are configured to receive audio data generated by members of Game Logic **210** and encode this audio data to audio packets. The encoded data packets are then sent to Clients **110** using I/O Device **145** and Network **115**. Audio DEPs **270** are also optionally configured to place an audio overlay on the audio data received from Game Logic **210**. Different Audio DEPs **270** are individually labeled **270A**, **270B**, **270C**, etc.

As with Video DEPs **245**, Audio DEPs **270** are typically allocated as needed and assigned to process audio data from a specific member of Game Logic **210**. This audio data may be based on events within a video game. For example, actions taken by one or more players of the video game may result in audio data according to game rules (e.g., falling in water makes a splashing sound). Allocation of Audio DEP **270A** is typically responsive to the same factors and performed in the same manner as allocation of Video DEP **245A**. For example, the elements included in Audio DEP **270A** may be responsive to needs of one of Clients **110**. Audio DEP **270** may receive audio data from Shared Memory **240** and/or directly from one of Virtual Audio Drivers **285**.

Audio DEP **270A** optionally includes a Grabber **275** configured to receive audio data from one of Virtual Audio Drivers **285**, to receive audio data from one or more sources external to Video Server System **120**, and/or to get audio data

18

from Shared Memory **240**. For example, in some embodiments Grabber **275** is configured to retrieve data received from sources external to Video Source **130**, such as one or more of Clients **110**, over the internet. The audio received, from any of these sources, can include voice audio, music, sound effects, sound tracks, alerts, and/or the like. For example, the voice of a game player using Client **110A** may be received from Client **110A**, or the game player may provide music to be heard by a game player using another of Clients **110**. Audio data may be received, as audio streams, from more than one game player in parallel.

Audio DEP **270A** optionally further includes Overlay Logic **295**. Overlay Logic **295** is configured to place one or more audio overlays on audio data generated by a member of Game Logic **210**. The audio overlays can include audio, music, sound effects, sound tracks, alerts, and/or the like. For example, in some embodiments, an audio overlay includes voice data received as data packets from one or more game players over the internet. These data packets are optionally decoded by Overlay Logic **295** prior to being included in the overlay. Placement of overlays on the audio data generated by a member of Game Logic **210** optionally includes scaling of the audio data within the overlay and addition of this data to the generated audio data. All or part of an audio overlay can be received from a source external to Audio Source **130**. For example, in some embodiments the overlay includes audio data received at I/O Device **145** via the internet and Overlay Logic **290** includes audio processing logic configured to generate an audio overlay from this audio data. When referring to audio overlays, the identities of which data is the overlay and which is the data on which the overlay is being applied may be arbitrary as, in some embodiments, audio data does not have a sense of order that can be possible in video data.

Audio DEP **270A** further includes an Audio Encoder **280A**. Audio Encoder **280A** is configured to encode the audio data generated by a member of Game Logic **210** and the audio overlay into an audio packet according to an audio codec. The encoding is typically performed after the audio overlay(s) have been applied to the audio data.

Grabber **275**, Audio Encoder **280A** and Overlay Logic **295** include hardware, firmware, and/or software stored on a computer readable medium. Audio DEP **270B** and Audio DEP **270C** are alternative embodiments of Audio DEP **270A**.

In some embodiments, Controller **265** is configured to provision instances of Game Logic **210A**, Virtual Audio Driver **285** and Virtual 3D Video Driver **220** in response to a request to play a game. For example, if a request for a specific video game is received from Client **110A**, Controller **265** may retrieve Game Logic **210A**, Virtual Audio Driver **285** and Virtual 3D Video Driver **220** from Storage **155** and place these in working memory. Game Logic **210A** may be placed in communication with both Client **110A** and Game Server **125**. Game Logic **210A** is then executed using Processor **150** in order to play the video game.

In some embodiments, Controller **265** is configured to allocate instances of Video DEP **245** and/or Virtual Audio Driver **285** in response to a request to play a game. For example, Controller **265** may be configured to first determine or receive the requirements for Video DEP **245A** and Audio DEP **270A** and then provision computing instructions that satisfy those requirements. Video DEP **245A** and Audio DEP **270A** are then associated with one of Game Logic **210** and one of Virtual 3D Video Driver **220**. The provisioned Video DEP **245A** is optionally configured to retrieve raw video data from a specific area within Shared Memory **240**.

In some embodiments the requirements for Video DEP **245A** and/or Audio DEP **270A** are determined by querying a

19

user of Client 110A. For example, Controller 265 may be configured to cause text messages to appear on Client 110A, the messages requesting that a user of Client 110A enter characteristics of a video stream desired by a user. These characteristics can include a connection type, a frame size, an encoding scheme, a frame rate, a color depth, and/or the like. The entered characteristics are conveyed back to Controller 265, which then used these characteristics to select elements to include in Video DEP 245A and/or Audio DEP 270A when they are provisioned. For example, if the user requests a specific encoding scheme, then an Encoder 255 configured to use the requested encoding scheme is selected and included in Video DEP 245A.

In some embodiments the requirements for Video DEP 245A and/or Audio DEP 270A are determined automatically. For example, information gathered using Client Qualifier 160 in an Execute Agent Step 310 (See FIG. 3) can be used to determine the requirements for Video DEP 245A. In this example, if Client 110A includes only a Microsoft Silverlight® decoder and an Adobe Flash® decoder, then at least one of the corresponding encoder configurations or variants thereof is required in Video DEP 245A. Characteristics of the video stream to be generated using Video DEP 245A, such as frame size, color depth and frame rate, can also be determined by Controller 265 in this way.

The determination of requirements for Video DEP 245A and/or Audio DEP 270 may include both the use of characteristics provided by a user and characteristics automatically determined. In some embodiments, the possible characteristics are limited by a user's account type, payment and/or subscription. For example, a user may pay to receive a higher quality video stream.

Controller 265 may also manage sharing of Shared Memory 240. For example, Controller 265 may configure Video DEP 245A and/or Audio DEP 270A to use specific memory locations within Shared Memory 240 and/or Controller 265 may be configured to configured Video DEP 245A and/or Audio DEP 270A to access Shared Memory 240 in response to certain timing signals. The timing signals being timed for synchronization of access to Shared Memory 240 by different Video DEPs 245.

FIG. 3 illustrates methods of receiving a qualified request for a video stream from a client such as Client 110A, according to various embodiments of the invention. A qualified request, as the term is used herein, is a request that is received from a client, where the client has been qualified to receive the video stream. The qualification means, for example, that the client meets a minimum level of decoding capability and/or that the communication channel between the client and Video Server System 120 meets minimum bandwidth and/or reliability requirements. In alternative embodiments, the request for a video stream is received from an unqualified client and the client is then qualified before the video stream is provided to the client.

In a Receive Request Step 300, a request is received to display an object on Client 110A. This object includes HTML, XML, text, an image, and/or the like, and is configured to indicate that a video stream is available for viewing on Client 110A. For example, the object may be a banner advertisement or some other type of link that indicates that a user can play a video game using Client 110A by clicking/selecting/cursor over the object. The request may be received as the result of the user loading a web page including a reference to the object. This reference can include a universal resource locator, HTML, XML, and/or the like. For example, in some embodiments a first user chooses to include a link and image of a video game they play on their web page within a social

20

networking website. When a second user visits this page information about the video game is displayed to the second user. As a result of computing instructions included within the web page, Video Server System 120 or some other part of Game System 100 receives a request to display text (an object) indicating that the second user can play the video game in a streaming mode by clicking on the object within the displayed information.

The identity of the Video Server System 120 or Game System 100 that receives the request in Receive Request Step 300 is optionally dependant on the nature of the request. For example, a request may include a URL that references a particular game publisher and this request is directed to an embodiment of Game System 100 controlled or licensed to this game publisher. The request may come from a specific web page and the identity of that web page may be used to determine which entity receives the request.

In a Provide Agent Step 305, a software agent is provided to Client 110A from Video Server System 120. The software agent is configured for discovering capabilities of Client 110A and/or capabilities of the communication channel between Client 110A and Video Server System 120. The software agent is configured to be executed on Client 110A using an electronic processor. For example, the software agent may be configured to examine a registry on Client 110A, identify video decoders installed on Client 110A, identify characteristics of a display of Client 110A, identify an electronic processor of Client 110A, identify firewalls installed on Client 110A, identify other software or hardware on Client 110A, and/or the like. In some embodiments Client 110A is configured to test the communication channel between Client 110A and Video Server System 120 by sending and receiving data between Client 110A and Video Server System 120. This data may include test video frames. In some embodiments more than one path between Client 110A and Video Server System 120 is tested. The characteristics of the communication channels tested may include, ping time, lag, jitter, bandwidth, reliability, and/or the like.

In some embodiments the software agent is configured to test communication channels between Client 110A and several different instances of Video Server System 120. In these embodiments, the software agent is capable of determining which of the communication channels is most capable of delivering a video stream to Client 110A from a member of the instances of Video Server System 120. As such, the software agent can determine which instance of Video Server System 120 is preferred for supplying the video stream.

In some embodiments the software agent includes computing instructions embedded in a web page. In some embodiments the software agent provided in Provide Agent Step 305 includes a browser plug-in. Provide Agent Step 305 is optional in embodiments where the software agent is already installed on Client 110A. In some embodiments, the software agent is a standalone program.

In an Execute Agent Step 310 the software agent provided in Provide Agent Step 305 is executed. The execution of the software agent may be automatic on delivery of the software agent to Client 110A, or the user loading the web page. The execution of the software is optionally manually initiated by the user. The software agent is executed using an electronic processor of Client 110A. The execution of the software agent is optionally initiated by a browser. The execution includes testing of communication channels and/or discovery of client capabilities as discussed herein.

In a Receive Capabilities Step 315 the information discovered by the software agent is received by Client Qualifier 160. The information typically includes the capabilities and fea-

21

tures of Client 110A as well as characteristics of one or more communication channel between Client 110A and one or more Video Server System 120. For example, after the execution of the software agent on Client 110A the information collected by the agent may be communicated to I/O Device 145 via Network 115.

In an Approve Client Step 320, Client 110A is approved by Client Qualifier 160 for the purpose of receiving a video stream from Video Source 130. Approve Client Step 320 includes comparing the information received in Receive Capabilities Step 315 with a set of minimal requirements for the delivery of streaming video to Client 110A. These requirements can include minimum capabilities of Client 110A and/or minimum requirements for the connection between Client 110A and Video Server System 120. The approval is made only if the minimum(s) are met or acceded. If the minimum(s) are not met or acceded then Client 110A is not approved for receiving the video stream from Video Source 130. In some embodiments the minimal requirements for delivery of streaming video are selected so as to assure that the video is presented to a user with at least a desired level of quality. The minimal requirements are optionally dependent on the identity of the video game associated with the request or a publisher thereof.

Approve Client Step 320 optionally includes identifying which of a plurality of alternative instances of Video Server Systems 120 is best suited to serve the streaming video to Client 110A. For example, the instance with the best communication channel between Client 110A and the Video Server Systems 120 may be selected. Criteria for determining which of Video Server Systems 120 is best suited to serve the streaming video are discussed elsewhere herein.

In some embodiments Client Qualifier 160 is configured to correct deficiencies in the capabilities of Client 110A. For example, Client Qualifier 160 can be configured to suggest to a user of Client 110A that a specific decoder be installed, or configured to automatically install a decoder. Client Qualifier 160 can also be configured to suggest to the user that hardware components of Client 110A should be upgraded.

Once a Client 110A is approved the approval may be stored, for example, on Client 110A or at Client Qualifier 160. If the approval is stored, then Client 110A is optionally pre-approved the next time Receive Request Step 300 occurs and Steps 305-320 are optional. When Receive Request Step 300 occurs again, the stored approval is retrieved and used to determine that Client 110A has already been approved. The delivery of different video streams may require different capabilities of Client 110A and/or of the connection between Client 110A and Video Source 130. Therefore, a stored approval optionally includes information characterizing the type of video stream approved and/or the capabilities received in Receive Capabilities Step 315. In some embodiments, stored information is used to qualify the capabilities of Client 110A, but the connection between Client 110A and Video Server System 120 is retested.

In a Display Object Step 325, the object requested in Receive Request Step 300 is displayed to the user of Client 110A. In some embodiments, the object is only displayed after Client 110A has been successfully approved in Approve Client Step 320. In other embodiments the object is displayed but not fully functional until Client 110A has been successfully approved in Approve Client Step 320.

For example, in some embodiments, the object requested in Receive Request Step 300 includes a banner advertisement. The request is made by code embedded in a web page. When the web page is loaded in a browser of Client 110A a request is sent to Client Qualifier 160 and Client Qualifier 160

22

responds by performing Steps 305-320 as needed. Following approval of Client 110A the banner advertisement is delivered to Client 110A and displayed on the web page in Display Object Step 325.

Alternatively, the banner advertisement is displayed on the web page prior to Approve Client Step 320, however, the functionality to which the banner advertisement is linked to is not fully activated. E.g., unlike a qualified banner advertisement, clicking on the banner advertisement will not result in delivery of a video stream. Delivery of the video stream is still only performed after Approve Client Step 320 is successfully completed.

In alternative embodiments Display Object Step 325 is performed prior to Execute Agent Step 310 and without Approve Client Step 320. In these embodiments, while the object is displayed it may not be fully functional and/or the actual delivery of the requested video stream is delayed until after Client 110A is approved in Approve Client Step 320. Thus, the video stream is still not provided to Client 110A until Client 110A is qualified. This assures that the video stream will be seen by a user of Client 110A at at least a minimum level of quality.

In a Receive Video Request Step 330, a request for streaming video is received by Video Server System 120 from Client 110A. This request and the streaming video are optionally communicated via Network 115. In some embodiments, the request is received as a result of a user clicking on, cursor over or selecting the object displayed in Display Object Step 325. In typical embodiments the streaming video is part of a video game played by the user of Client 110A. As such, the content of the streaming video is responsive to game rules and game commands entered by the user, and optionally users of other Clients 110. The request for streaming video may be in the form of a request by the user to play a video game in a streaming mode.

In some embodiments, Video Server System 120 is configured such that the request received in Receive Video Request Step 330 causes game and video related resources to be allocated within Video Source 130. FIG. 4 illustrates methods of allocating these resources within a Video Source 130, according to various embodiments of the invention. These methods include provisioning of Game Logic 210A and Virtual 3D Video Driver 220, and/or dynamic provisioning of Video DEP 245A. The methods illustrated by FIG. 4 are optionally performed using a system manager and a code source discussed elsewhere herein.

In a Determine Needs Step 410, the needs of Client 110A are determined. These needs may include a desired frame rate, frame size, color depth, encoding, minimum lag time, other needs discussed herein, and/or the like. The needs of Client 110A may be dependent on a particular video game played, for example some video games require a lower lag time between the time a game player makes a game command and the video the game player sees reflects the results of the game command, relative to other video games.

In some embodiments, the needs of Client 110A are dependent on selections made by a game player. For example, a player may select a specific frame size or set the size of a browser window that a frame should fill. In some embodiments, the needs of Client 110A are selected to simulate video that would be produced by specific hardware. For example, a game player may select to see video as would be produced by a specific nVidia or a specific AMD graphics card. In this case, the needs are typically specified to reflect the maximum performance of the selected device.

In some embodiments the needs of Client 110A are determined by accessing data received in Receive Capabilities

23

Step 315 (FIG. 3) and optionally stored in Storage 155. This data is described elsewhere herein. The needs of Client 110A are then selected so as to match the capabilities of Client 110A. For example, if Client 110A includes an Adobe Flash decoder then one of the needs may be a video stream encoded in Adobe Flash. If communication to Client 110A has a limited bandwidth, one of the needs may be a limit on the frame size that can be transmitted to Client 110A.

In some embodiments, Determine Needs Step 410 includes providing an agent to Client 110A, executing that agent and receiving capabilities of Client 110A. This approach may include performing embodiments of Provide Agent Step 305, Execute Agent Step 310 and Receive Capabilities Step 315 (FIG. 3).

In some embodiments, Determine Needs Step 410 includes reading a cookie, IP address, or other identifying information on Client 110A. This identifying information is then used to look up previously stored needs for Client 110A. For example, if the needs of Client 110A were previously determined, these needs can be stored in a database in, for example, Storage 155 or some location remote from Video Server System 120. The identifying information can then be used to retrieve these stored needs. Determine Needs Step 410 optionally includes both retrieving stored information and performing some tests on Client 110A. For example, information regarding software available on Client 110A may be stored while the connection between Client 110A and Video Source 130 may be tested again.

In some embodiments, Determine Needs Step 410 is dependent on minimum requirements set by a provider of a video game. For example, a game publisher may specify that a video stream be provided at at least a minimum resolution and/or quality of service. This enables a game publisher to assure that a game demo meets minimum standards.

In a Determine Video Source Step 415 a Video Source 130 is selected for providing streaming game video to Client 110A. The selected Video Source 130 is typically selected from among a plurality of Video Sources 130. This plurality may include Video Sources 130 located at the same location or geographically distributed. The selection is based, in part, on the ability of the Video Source 130 to meet the needs of Client 110A as determined in Determine Needs Step 410. For example, if a first Video Source 130 can provide streaming video with a lower lag time relative to a second Video Source 130, the first Video Source 130 may be more likely to be selected. Further details relating to the selection of Video Sources 130 are discussed elsewhere herein. For example, in relation to FIG. 7.

In a Select Game Logic Step 420 an instance of Game Logic 210 is selected for inclusion in the Video Source 130 selected in Determine Video Source Step 415. The Game Logic 210 is typically selected based on which video game the game player wishes to play. The Game Logic 210 may also be selected based on the needs of Client 110A. For example, if Client 110A includes a touch screen interface, a version of the video game (Game Logic 210) configured for use with a touch screen interface may be selected over a version configured for a keyboard interface. In some embodiments, a demo or introductory version of a video game may be selected over a full version responsive to a user's account status, e.g., a user making a payment. An instance of Game Logic 210 is optionally selected based on a need for a specific video frame aspect ratio. Different instances of Game Logic 210 may be configured for the generation of video frames having different aspect ratios.

In a Provision Game Logic Step 425 the Game Logic 210 selected in Select Game Logic Step 420 is provisioned onto

24

the Video Source 130 selected in Determine Video Source Step 415. Provision Game Logic Step 425 typically includes copying computing instructions included within Game Logic 210 from Storage 155 or from a code source external to Video Server System 120 to working memory of Video Source 130. Provision Game Logic Step 425 is optional where the selected Game Logic 210 is already provisioned on the selected Video Source 130.

In a Provision (Virtual 3D Video) Driver Step 430 in instance of Virtual 3D Video Driver 220 is provisioned to the instance of Video Source 130 selected in Determine Video Source Step 415. The Virtual 3D Video Driver 220 is typically retrieved from Storage 155 and copied to working memory of Video Source 130. The Game Logic 210 provisioned in Provision Game Logic Step 425 and the Virtual 3D Video Driver 220 are also linked together (e.g., a memory address of Virtual 3D Video Driver 220 is provided to Game Logic 210) such that the game environment output of Game Logic 210 is directed toward the Virtual 3D Video Driver 220.

In an optional Select Preprocessor Step 435 one or more Preprocessor 250 are selected based on the needs of the Client 110 as determined in Determine Needs Step 410. The selection may also be based on the type of output of the Graphics Processing Unit 235 included within the selected Video Source 130. Select Preprocessor Step 435 is optional wherein raw graphic provided by Graphics Processing Unit 235 is directly compatible with an encoder.

In a Select Encoder(s) Step 440, one or more video Encoders 255 are selected for inclusion in Video Source 130. One of this Encoders 255 is typically a preferred video encoder based on the video needs of the client. Others of these encoders may be backup Encoders 255 in case there is a need to change encoding during delivery of a video game. For example, a preferred Encoders 255 may produce a higher quality video stream than a backup Encoders 255 but the backup Encoders 255 may require less bandwidth for delivery of video to Client 110A. Typically, the selection of Encoders 255 will be based on at least the video decoders available on Client 110.

In an optional Select Post Processor(s) Step 445, one or more Post Processors 260 are selected for inclusion in Video Source 130. The selection is dependent on what further processing/packaging the output of the Encoders 255 selected in Select Encoder(s) Step 440 require.

In a Provision DEP Step 450 a Dynamic Encoding Pipeline (Video DEP 245A) is provisioned in Video Source 130 using the Preprocessor 250, Encoders 255 and/or Post Processors 260 selected in the previous steps. Provision DEP Step 450 includes linking these elements together such that they can pass access to video data received from Shared Memory 240 and/or Graphics Processing Unit 235. For example, in some embodiments memory addresses are provided to these elements so they can retrieve video data from Shared Memory 240 and then operating on that video data in a Heap within Video DEP 245A. This step may be performed under the control of Controller 265.

Provision DEP Step 450 typically includes copying copies of the provisioned elements from Storage 155 or from a code source discussed elsewhere herein. If more than one Encoder 255 is selected in Select Encoder(s) Step 440, just the preferred Encoder 255 or alternatively more than one Encoder 255 may be initially provisioned in Provision DEP Step 450. If only one encoder is initially provisioned additional encoders may be provisioned at a later time. Video DEP 245A is optionally configured in a tree structure. For example, one Preprocessor 250 may provide video frames to one or more Encoders 255 and each of Encoders 255 may provide video frames to one or more Post Processors 260.

25

In some embodiments, the creation of Video DEP 245A in Provision DEP Step 450 includes receiving a command. This command includes parameters specifying the elements to include in Video DEP 245A. Typically, a Video DEP 245 is not allocated and provisioned until needed by a specific Client 110 at which time it is provisioned for that client.

FIG. 5 illustrates methods of serving multiple video games, according to various embodiments of the invention. In these methods Video Source 130 supplies multiple video streams to multiple destinations. These video streams are based on game environments generated by Game Logic 210 and are rendered using shared Graphics Processing Unit 235. The steps illustrated in FIG. 5 may be performed in parallel to generate the multiple video streams. For example, Game Logic 210A, Game Logic 210B and Game Logic 210C may each be generating game environments for rendering and the resulting video streams can be sent to different game Clients 110. The operation of this Game Logic 210 may or may not be coordinated. For example, in some embodiments the Game Logic 210 is synchronized so as to produce game environments at different times so that Graphics Processing Unit 235 can render each game environment as quickly as possible.

In a Receive State Step 510, one or more game states are received by Game Logic 210. Each game state is directed to a different member of Game Logic 210. For example, if Video Source 130 is configured to execute a client (Game Logic 210) for World of Warcraft® using Game Logic 210A and a client for Grand Theft Auto® using Game Logic 210B, then game states for each of these games are received and directed to the respective Game Logic 210. Each of the game states are representative of a state of the respective video games.

The game states can be received from Game Server 125 or from a storage location within Video Source 130. In some embodiments, the receive game states are the result of game commands received by Game Logic 210A and Game Logic 210B and passed on to one or more Game Server 125, and/or the result of game commands received by Game Server 125 from other game players. In various embodiments a received game state represents a saved game or a game state updated locally to Video Source 130.

In a Determine 3D Data Step 515, the game states received by instances of Game Logic 210 are each used to determine a separate first three-dimensional game environment based on the game state and a player point of view. The game environments are determined using rules of the game as defined by each Game Logic 210. Each game environment is passed to a Virtual 3D Video Driver 220.

In a Render Step 520, the one or more three-dimensional game environments generated in Determine 3D Data Step 515 are rendered to video data. This rendering is performed using shared Graphics Processing Unit 235. In some embodiments, the rendering of each three-dimensional game environments are done in parallel. For example, different game environments may be processed by different parts of Graphics Processing Unit 235 at the same time, or Graphics Processing Unit 235 may be configured to process different game environments in different time slices. These time slices may or may not be long enough to generate a single video frame.

In some embodiments, game environments are provided to Graphics Processing Unit 235 in a synchronized manner such that Graphics Processing Unit 235 can process one or a limited number of game environments at the same time. The synchronization can be accomplished by communication between members of Game Logic 210, or by control of Virtual 3D Video Drivers 220. For example, Virtual 3D Video Drivers 220 may be configured to report back completed processing of game environments to each of Game Logic 210

26

in a synchronized manner such that the next game environment is provided from each of Game Logic 210 at the best time. This synchronization is optionally managed by Controller 265. The result of Render Step 520 is typically a separate frame of raw video data for each receive game environment.

In a Transfer Step 525, the raw video data generated by Graphics Processing Unit 235 is transfer to Shared Memory 240. This transfer is optionally performed under the control of Virtual 3D Video Drivers 220. The raw video data is optionally transferred to specific memory locations within Shared Memory 240 assigned to specific Video DEP 245. Alternatively, the raw video data is transferred at specific times assigned to specific Video DEPs 245. For example, each of Virtual 3D Video Drivers 220 may be configured to control access by an assigned member of Video DEPs 245 to Shared Memory 240 such that the member of Video DEPs 245 receives the correct video stream.

In an optional Preprocess Step 530, the raw video data generated by Graphics Processing Unit 235 and received by Video DEP 245A is preprocessed using Preprocessor 250. Preprocess Step 530 typically includes the conversion of the raw video data from RGB data to YUV data. Preprocess Step 530 is optional in embodiments wherein the raw video data generated using Graphics Processing Unit is compatible with Encoder(s) 255. For example, if Graphics Processing Unit is configured to generate YUV data directly. Multiple instances of Preprocess Step 530 are optionally performed in parallel within different Video DEP 245 on different video streams generated by Graphics Processing Unit 235. Preprocess Step 530 optionally further comprises changing a frame size of the raw video data.

In an Encode Step 535, one or more Encoders 255, e.g., the preferred Encoder 255, are used to encode the video data generated by Graphics Processing Unit 235 and optionally processed using Preprocessor 250. The encoding is to one or more specific codec. In Encode Step 535 video data generated using Graphics Processing Unit 235 is optionally encoded to more than one codec or to more than one frame size, color depth, or other video characteristic. For example, Encoder 255A may be configured to encode according to one codec at a first frame size while Encoder 255B is configured to encode according to a second codec at a different frame size.

In an optional Package Step 540, the encoded video frame generated in Encode Step 535 is packaged for communication to a Client 110, if needed.

In a Provide Step 545, the video frame encoded in Encode Step 535 and optionally packaged in Package Step 540 is provided to one or more Client 110 over Network 115 as part of a video stream. This video frame may also be provided to other destinations. For example, a copy of the video stream may be delivered to a website, a storage location, a social networking site, an advertisement, and/or the like.

The steps illustrated in FIG. 5 are optionally performed in different orders and may be repeated with variations. For example, Encoder 255A may be used in one instance of Encode Step 535 while Encoder 255B is used in another instance of Encode Step 535. This change may occur during delivery of a video stream. Encoder 255A and Encoder 255B optionally communicate in order to maintain a sequence numbering of the video frames within the video stream. Reasons for changing encoding or other characteristics of the video stream are discussed elsewhere herein.

FIG. 6 illustrates various embodiments of Game System 100 including more than one Video Server Systems 120. These Video Server Systems 120 are individually identified as 120A, 120B, 120C, etc. Game System 100 can include at least 1, 2, 3, 5, 10, 25, 50, 100 or 500 instances of Video Server

27

System 120. These embodiments of Game System 100 further include a System Manager 615 configured to manage the use and operation of Video Server Systems 120. These embodiments of Game System 100 optionally also include a Code Source 625 configured to provide game code to Clients 110 at the same time these Clients 110 are receiving streaming video.

The members of the Video Server Systems 120 illustrated in FIG. 6 are typically geographically dispersed. For example, Video Server System 120A may be located in San Jose Calif., Video Server System 120B may be located in Tokyo, and Video Server system 120C can be located in New York City. The geographic distribution of Video Server Systems 120 allows Clients 110 to receive video streams over a shorter geographic distance that would occur if all of Video Server Systems 120 were located in the same place. To the extent that geographic distance is related to the time required to communicate between one of Clients 110 and one of Video Server Systems 120, reducing the geographic distance reduces the latency of this communication.

Each of Video Server Systems 120 are optionally configured to provide multiple video streams to multiple Clients 110. These video streams can arise from different video games and/or multiple instances of the same video game. Different members of Video Server Systems 120 may be configured to provide streaming video from different sets of video games. For example, Video Server System 120B may be configured to generate video streams arising from 20 instances of World of Warcraft® and 15 instance of Mortal Combat®, while Video Server System 120B are configured to generate video streams arising from 20 instance of World of Warcraft and 25 instances of Spore®.

In various embodiments, Video Server Systems 120 are configured to communicate with Clients 110 and optionally one or more Game Servers 125 via Network 115. Typically Network 115 includes the internet or other public network that carries network traffic other than that related to the output of Video Server Systems 120. Members of Video Server Systems 120 optionally communicate through a Proxy Server 605 and a Dedicated Connection 610. Proxy Server 605 is a network server that is configured to appear to be one of Video Server Systems 120 from the point of view of Clients 110, and vice versa. The purpose of Proxy Server 605 is to forward received messages, e.g., from Client 110A to Video Server System 120C and from Video Server System 120C to Client 110A. The advantage of using Proxy Server is that part of the communication can occur over Dedicated Connection 610. Dedicated Connection 610 can have a high bandwidth dedicated to this communication. In one example, Proxy Server 605 and Client 110A are both located in Amsterdam. Dedicated Connection 610 is configured to enable high bandwidth communication between Proxy Server 605 and Video Server System 120C in New York City. Communication between Video server System 120C and Client 110A occurs by first communicating between Video Server System 120C and Proxy Server 605 via Dedicated Connection 610, and then by communicating between Client 110A and Proxy Server 605 via Network 115. The use of Proxy Server 605 can allow a game player in New York to play on the same member of Video Server Systems 120 as a game player in Amsterdam.

System Manager 615 includes hardware, firmware, and/or software stored on a computer readable medium. For example, System Manager 615 can include an Internet server having storage, an electronic processor, and software configured to operate the electronic processor so as to perform actions discussed herein. System Manager 615 may be located at a same location as one of Video Server Systems 120

28

or at a separate location. Typically, System Manager 615 includes at least an embodiment of Client Qualifier 160 and a Monitor 620.

The embodiment of Client Qualifier 160 within System Manager 615 is configured to qualify Clients 110 based on the capabilities of the Clients 110 and the connections between Clients 110 and Video Server Systems 120. When multiple Video Server Systems 120 are available to provide video to Client 110A, Client Qualifier 160 optionally includes computer instructions configured to first determine a geographic location of Client 110A based on mapping of the IP (Internet protocol) address. This mapping is then used to identify one or more of Video Server Systems 120 in geographic proximity to Client 110A. Following the identification of nearby sources of streaming video, the communication channels between the identified members of Video Server Systems 120 are tested using further computing instructions of Client Qualifier 160. This testing is further described elsewhere herein. Client Qualifier 160 is further configured to use the results of these tests to determine which of the Video Server Systems 120 is best qualified to deliver streaming video to Client 110A.

Monitor 620 is configured to manage Video Server Systems 120. For example, Monitor 620 can track the usage and capacity of each of Video Server Systems 120. This can include the geographic location of each of Video Server Systems 120, the number of video streams provided by each, the total amount of video data provided by each, the percentage of GPU and CPU cycles that are used, the amount of storage available, temperature, fan speed, hard drive and memory usage, network card usage, and/or the like. Monitor 620 can also track which games are provisioned on which of Video Server Systems 120 and how much each game is played. In some embodiments, Monitor 620 is configured to provision members of Video Server Systems 120 with Game Logic 210, Encoders 255, and/or other software needed to provide video streams associated with specific games. Monitor 620 may also be configured to manage assignment of a specific member of Video Server Systems 120 to provided streaming video to a specific member of Clients 110, based on testing performed by Client Qualifier 160. This assignment can include reassignment from one member to another member of Video Server Systems 120. Reassignment is desirable under conditions discussed elsewhere herein.

Code Source 625 is configured to provide computing instructions to Video Server Systems 120, Video Source 130, Clients 110, and/or other elements of Game System 100. For example, Code Source 625 may be configured to provide Game Logic 210, Preprocessor 250, Encoder 255, Post Processor 260, Operating System 215, a software agent for testing the capabilities of Client 110, and/or the like. Code Source 625 may also be configured to provide game code to Clients 110 as a background process as streaming video is being provided to these Clients 110. In typical embodiments, Code Source 625 includes a storage device such as a hard drive, a file server, a flash memory, an optical storage, or other device configured to store computing instructions. Code Source 625 is optionally configured to communicate with other parts of Game System 100 via Network 115.

FIG. 7 illustrates methods of providing a video stream to a qualified client, according to various embodiments of the invention. In this method one of Clients 110, e.g., Client 110A is qualified to receive video, based on the qualification a user of Client 110A is given an opportunity to request the video, and upon receiving the request streaming video is provided. The method is optionally performed using Game System 100. For example, the steps illustrated in FIG. 7 can be managed by

29

System Manager **615** and/or Controller **265**. While the steps illustrate in FIG. 7 are explained with reference to delivering a video stream to a single instance of Client **110**. In alternative embodiments, the video stream may be delivered to multiple Clients **110**. In these embodiments, various steps such as tests may be performed with respect to each of the multiple Clients **110**. The steps illustrated in FIG. 7 may also be adapted to the delivery of audio streams.

The methods illustrated in FIG. 7 start with Receive Request Step **300** described elsewhere herein, e.g., with respect to FIG. 3. In this step a request is received to display an object on Client **110A**. This object is configured to give a game player an opportunity to request to receive a streaming video as part of playing a video game. The request could be for a specific video game, any of a class of video games, a “game of the day,” a game undergoing a sales promotion, the identity of a web page on which the object is displayed, or any video game. The request is optionally received at System Manager **615** or one of Video Server Systems **120**.

In an optional Check Cookie Step **705**, Client **110A** is examined for a browser cookie, an IP address, and/or other identifying information. If this information is found it is used to retrieve previously determined information regarding Client **110A**. For example, if characteristics and needs of Client **110A** have already been determined and stored, the identifying information is used to retrieve the stored characteristics and needs. The information may be stored on System Manager **615** or one of Video Server Systems **120**. The retrieval of previously stored information may eliminate the need to perform some or all tests in order to qualify Client **110A** to receive the object requested in Receive Request Step **300**.

In an optional Provide (Web) Page Step **710** a web page is provided to Client **110A**. This page may or may not have the requested object included therein depending on the particular embodiment of the invention. If the requested object is included it is typically not fully functional.

In an Identify Servers Step **715** one or more Video Server Systems **120** are identified as potential candidates for providing streaming video to Client **110A**. The identification of these servers may be based on their proximity to Client **110A**. For example, an IP address of Client **110A** may be mapped in order to determine its approximate location. The identification of these servers may also be based on an initial ping test between the various Video Server Systems **120** and Client **110A**. The identification of these servers may also be based on the available capacity of Video Server Systems **120**. For example, an instance of Video Server Systems **120** may not currently have capacity to deliver the streaming video to Client **110A**. In this case this instance of Video Server Systems **120** is not considered as a candidate. It may occur that no Video Server Systems **120** are identified as candidates for providing streaming video to Client **110A**. In this case, Client **110A** cannot be qualified to receive the streaming video and the method is halted.

In a Test Client Step **720**, Client **110A** and the connections between Client **110A** and the various Video Server Systems **120** identified as candidates in Identify Servers Step **715** are tested. Test Client Step **720** can include Provide Agent Step **305**, Execute Agent Step **310** and Receive Capabilities Step **315**, as described elsewhere herein. The tests made in Test Client Step **720** are optionally performed using Java® code. This code can be included in a browser plug in or embedded in a web page. For example, in some embodiments, the testing is performed just using Java code embedded in a web page.

The testing performed in Test Client Step **720** can include: latency tests such as the time required to communicate between Client **110A** and each of the candidate Video Server

30

Systems **120**; bandwidth tests based on, for example, HTTP or UDP; tests of the CPU on Client **110A**, tests configured to determine the availability of plug-ins on Client **110A** (plug-ins can include testing agents, decoders, or the like); tests configured to determine the operating system running on Client **110A**; tests configured to determine the types of devices included in or attached to Client **110A**, such as display devices or graphics processing units; tests configured to determine the size and aspect ratio of an area on Client **110A** on which streaming video is to be displayed; tests configured to determine a maximum packet size that can be transmitted to Client **110A** without being split (e.g., a maximum transmission unit); tests configured to determine if a Wifi connection is included in the communication channels between Client **110A** and each of the candidate Video Server Systems **120**; and tests configured to determine the reliability (e.g., jitter, packet loss and variation in bandwidth) between Client **110A** and each of the candidate Video Server Systems **120**.

The tests performed in Test Client Step **720** are optionally performed over multiple web page loads. These can occur over time, e.g., over days. In order to perform tests over time identification information, such as that identified in Check Cookie Step **705** is used to associate previous test data with new test data. For example, some reliability testing of communication channels between Client **110A** and candidate Video Server Systems **120** may occur each time a game player interacts with Game System **100** and the results of these tests can then be combined to create a statistically reliable estimate of the reliability. Tests that do not take significant time, e.g., ping tests, are typically performed on a single page load. While tests that take more time, e.g., bandwidth and reliability tests, are optionally performed over several page loads or several instances of Test Client Step **720**. Old test data is optionally allowed to expire such that testing reflects current conditions of Client **110A**. Test Client Step **720** is optionally performed, in part, after delivery of streaming video to Client **110A** has begun.

In a Determine Server Availability Step **725** the availability of the candidate Video Server Systems **120** identified in Identify Servers Step **715** to provide the streaming video are determined. This available may be dependent on resources available on each Video Serving System **120**. For example, whether sufficient Processor **150** or Graphics Processing Unit **235** capacity is available. If the request received in Receive Request Step **300** is directed toward a specific video game, then Determine Server Availability Step **725** can include confirming that the candidate Video Serving Systems **120** are configured to provide streaming video from that specific video game. An instance of Video Serving Systems **120** may be provisioned with required Game Logic **210** if needed.

In a Determine Game Availability Step **730** the availability of a specific video game is determined. If the request received in Receive Request Step **300** is directed toward a specific video game then the availability of that game is determined. Otherwise, the availability of several games may be determined until an available game is found. Game availability may also be dependent on licensing of game code from a publisher or a game promotion program.

The availability of a specific video game may be dependent on test results generated in Test Client Step **720**. For example, some video games require a greater amount of bandwidth and or shorter ping time than other video games. Optionally, if the tests results show that a computer game cannot be delivered with a minimum level of quality then that game is not available to Client **110A**, at this time. A game player using Client **110A**, a manager of Game System **100**, and/or a publisher of a video game may determine the minimum level of quality.

31

The minimum level of quality can include several parameters such as lag time, bandwidth, reliability, and frame size.

In a Select Server Step 735 one of the candidate Video Serving Systems 120 is selected to provide the video game, as streaming video, to Client 110A. This selection is based on the results of Test Client Step 720, Determine Server Availability Step 725, and/or Determine Game Availability Step 730. Typically the Video Serving System 120 best able to provide the streaming video to Client 110A is selected. Select Server Step 735 optionally includes temporarily reserving capacity (memory, CPU and GPU, etc.) on the selected Video Server System 120. This assures that the required resources are available if the video stream is actually delivered. Select Server Step 735 is an alternative embodiment of Determine Video Server Step 415 (FIG. 4).

In a Determine Needs Step 740 the needs of Client 110A are determined. Determine Needs Step 740 often includes merely retrieving information found in Test Client Step 720 and/or Determine Game Availability Step 730. For example, the need of a specific decoder can be determined based on the results of Test Client Step 720 and a minimum frame size can be determined based on the results of Determine Game Availability Step 730. Determine Needs Step 740 is an alternative embodiment of Determine Needs Step 410.

In a Provision Step 745 the selected Video Serving System 120 (and a Video Source 130 therein) are provisioned to provide the streaming video to Client 110A. Provision Step 745 optionally includes Select Game Logic Step 420, Provision Game Logic Step 425, Provision Driver Step 430, Select Preprocessor Step 435, Select Encoder Step 440, Select Post Processor Step 445, and/or Provision DEP Step 450, as discussed elsewhere herein. The provisioned logic and other computing instructions are optionally provisioned from Code Source 625.

In an Offer Game Step 750 the object requested in Receive Request Step 300 is provided to a game player on Client 110A. Offer Game Step 750 can include providing a pop-up browser window, changing an area of a web page already provided in Provide Page Step 710, sending an e-mail or text (SMS or MMS) message, sending a link, or direct streaming of game video. For example, in some embodiments a region of the web page provided in Provide Page Step 710 directly receives streaming game video with no further actions by the game player. The region within the web page can show, for example, streaming video of the starting line of Mario Karl, or the character selection screen or a starting zone of World of Warcraft. The game player can begin game play by just selecting the region and starting to provide game commands. Optionally, the region of the web page can show streaming video including the avatar of a friend of the game player playing the game and an avatar of the game player. This allows the game player to instantly join a game with a friend. Friends can be identified using relationships established on social networking sites, for example. In some embodiments, Offer Game Step 750 is delayed until capacity to provide the streaming video becomes available. Offer Game Step 750 is optionally automatically performed once the capacity is available.

Offer Game Step 750 is contingent on Client 110A being qualified to receive streaming video of the game and that the game be available. Offer Game Step 750 is typically automatic once these conditions have been satisfied. In some embodiments, Offer Game Step 750 includes adding functionality to a web page already displayed to a game player in Provide Page Step 710. In these embodiments the requested object is already displayed but not fully functional. The added

32

functionality makes it functional. The added functionality may be a proper link to a source of streaming video.

In a Deliver Game Step 755 the video stream is provided from the selected Video Server System 120 to Client 110A. Deliver Game Step 755 includes changing the contents of the video stream responsive to game commands received from the game player via Client 110A. Deliver Game Step 755 optionally includes providing one or more copies of the video stream to other destinations, e.g., to other IP addresses. These copies can have different characteristics than the video stream provided to Client 110A, but are based on the same game environment provided to Graphics Processing Unit 235. The other destinations include websites, storage locations, social networking sites, other Clients 110, and/or the like. The streaming video is optionally provided to all game players (and lurkers) within the same game room. A game room is a virtual room in which players can join, leave and play games as a group. A game room can include game players and optionally lurkers who observe but do not play the video game. Game rooms are useful when playing games designed for a group of players sitting around a game console and sharing a display device.

Deliver Game Step 755 optionally includes delivery of information other than the video stream to Client 110A. For example, Deliver Game Step 755 can include delivery of executable game code to Client 110A as a background process. This executable game code can allow eventual execution of the video game in a client side mode (where rendering occurs on Client 110A) rather than at Video Source 130.

In an optional Monitor Step 760 delivery of the streaming video to Client 110A is monitored. This monitoring can include making sure that the delivery of the streaming video meets a required quality of service. The monitoring can also include gathering further data characterizing the reliability and bandwidth of the communication channel between Client 110A and the selected Video Server System 120. In some embodiments, Monitor Step 760 includes monitoring for changes in the size of a browser window in which the streaming video is being presented on Client 110A. For example, if the game player changes the size of a browser window the monitoring may determine that a new frame size for the streaming video is warranted.

In an optional Add Game Player Step 765 a new game player and destination (e.g., Client 110) for the streaming video is added. This may occur, for example, when a new player enters a game room. The addition of a new game player may change the geographical distribution to which the streaming video is delivered. When this occurs, Identify Servers Step 715, Test Client Step 720, Determine Server Availability Step 725, Determine Game Availability Step 730 and/or Select Server Step 735 may be repeated.

In an optional Change Server Step 770 delivery of the streaming video from the member of Video Server Systems 120 selected in Select Server Step 735 is transferred to a different member of Video Server Systems 120. This transfer may include provisioning the appropriate resources on the new Video Server System 120, optionally synchronizing execution of Game Logic 210 executing on the two Video Server Systems 120, and transferring delivery of the video stream. Further details of this process are disclosed in U.S. patent application Ser. No. 12/790,955 filed May 31, 2010 and entitled "Game Execution Environments."

Using the methods illustrated in FIG. 7 the utilization of Video Server Systems 120 can be controlled. For example, as game players are not offered to play video games unless there is capacity available, spare capacity is not necessary. This allows for a high utilization rate (e.g., greater than 50, 75, 80,

90 or 95%) of Video Server Systems **120**. In various embodiments, more than 80, 60 or 50% of requests received in Receive Request Step **300** are denied, even with more than ten of Video Server System **120**. While the examples provided herein are primarily directed at the delivery of video, the systems and methods described herein can be applied to the generation of still images and the delivery of still images, and/or the delivery of audio frames and the delivery of audio streams. In some embodiments, both video and audio streams are generated and provided as “audio-visual” streams.

In some embodiments, one or more Video Server Systems **120** are operated in multiple modes. In a “game play” mode resources of Video Server Systems **120** are used to provide video games to video game players at times selected by the video game players, i.e., video games are provided on the game players demand. In this mode players can request games without regard to resource availability. At the same time, in a “demo” mode video game demos are provided at times controlled by the Video Server Systems **120**, i.e., game demos (that typically include playing the video game) are provided only when offered to a game player. Offers are only made when resources are available for play and, thus, a player can only request a game when resources are available. Resources of the Video Server Systems **120** are allocated to the demo mode so as to minimize spare resource capacity and thus maximize utilization. The number of demos offered is dynamically adjusted according to the resources that can be allocated to the demo mode. Using this approach one or several of Video Server System **120** can include the resources needed to support the maximum number of simultaneous users expected in the game play mode, and at times when less than the maximum number of simultaneous user are being supported, demos are offered to fill the access capacity. The demos can be of different or the same video games. Systems and methods of serving demos are further described in U.S. patent application “Dual-Mode Program Execution” and having Ser. No. 13/231,790. Balancing of the game play mode and demo mode is optionally coordinated using elements of System Manager **615** or “Game Delivery Systems **315**” as described in the U.S. patent application “Game Execution Environments” and having Ser. No. 12/790,955. The above U.S. patent applications are hereby incorporated herein by reference.

FIG. **8** illustrates overlays on game video (a video frame), according to various embodiments of the invention. Three Overlays **810A**, **810B** and **810C** (collectively identified as Overlays **810**) have been applied to the video generated by the video game. These Overlays **810** may be applied as a single overlay or as several separate overlays. Each of Overlays **810** represent input controls and are clickable objects associated with an underlying input command. For example, the “TAB” Overlay **810C** may be part of an input control that generates a set of screen coordinates when clicked, the screen coordinates being mapped (at Client **110A** or Video Server System **120**) to a “TAB” keystroke. In this example, the screen coordinates and the “TAB” keystroke are both examples of input commands. FIG. **8** further includes examples of Input Controls **815** that are not overlays, but are included in the video frame by Game Logic **210A**. These clickable Input Controls **815** are configured to generate different game commands. In some embodiments, an overlay includes an image of a cursor.

Each of the Overlays **810** represents a different game input control as indicated by the text therein. Overlay **810A** represents a carriage return, Overlay **810B** represents an escape key button, and Overlay **810C** represents a Tab key button. In typical embodiments, these Overlays **810** are associated with specific regions of the video frame and these specific regions

are mapped to the represented game input controls. For example, the region of the image occupied by Overlay **810C** is mapped to the “Tab” key such that clicking on or otherwise making a selection in this region is interpreted as a Tab key-stroke. Selection can be accomplished by using a pointing device such a joystick or mouse, or by touching this region on a touch sensitive screen.

FIG. **8** further includes an Overlay **820** representing a game menu, such as may be used to manage a game room or other aspect of game execution. Overlay **820** includes a plurality of active regions mapped to different game commands (inputs). For example, the area occupied by the word “Search” within Overlay **820** is mapped to a search command while the area occupied by the word “Join” is mapped to a command to join a game room.

FIG. **8** further includes an Overlay **830** including an image of a game player. Overlay **830** may be a static or real-time image of another game player. For example, if the video frame illustrated in FIG. **8** is to be sent to Client **110A** then Overlay **830** may include a real-time video frame of a player of Client **110B**. Overlay **830** may alternatively include a view of the video game from another player’s point of view, a view from a second point of view for a same player, or a view of some real-life event. In one example, Overlay **830** includes a video frame based on a second point of view of the same player whose first point of view was used to generate the underlying image. This second point of view can be generated using a second Video DEP **245**. The number, shape, size and characteristics of overlays that may be placed over video data can vary widely. Overlays may or may not be mapped to game commands. Overlays can be applied either server and/or client side. Some embodiments include overlays applied both server and client side.

The mapping of user input controls to game commands is not limited to applications in which Clients **110** include a touch screen or pointing device. FIGS. **9A**, **9B** and **9C** illustrate game control mappings between Buttons **910** of a Game Controller **920** and game commands, according to various embodiments of the invention. These game commands include attack actions, movement actions, changing views, menu display request items, cursor control commands, targeting commands, save actions, video setting menu items, print screen, and/or the like.

The mapping of user input controls to specific game input commands can occur at Clients **110** or at Video Server System **120**. For example, the iPad® from Apple Computer, Inc. is configured to convert touches to its screen, movements, and combinations of touches to specific commands and further have the capacity to remap these commands. The iPad and iPhone® are also configured to convert movements of the device itself to specific commands. However, other types of Client **110** may not have these capabilities, in which case the mapping can occur at Video Server System **120**.

Mapping of user input controls to game commands (and, more generally input commands to other input commands) is optionally performed using a look-up table or other type of input mapping. FIG. **10** illustrates part of an Input Map **1010** configured for mapping of user input controls to game commands, according to various embodiments of the invention. Typically, an Input Map **1010** for command mapping will include more entries than illustrated. The mappings within Input Map **1010** are optionally dependent on a state of the video game, as described further elsewhere herein. Further, mappings within Input Map **1010** are optionally dependent on characteristics of a game player’s account, e.g., a premium account may be allowed commands that a standard account is not.

35

The first mapping illustrated in Input Map 1010 is a direct character to character map of “D” to “K.” This type of mapping may be useful, for example, in games that are configured for several players to share a keyboard, each using their own subset of the keys. When such a game is played by geographically

disperse game players each player may use the same keys but one of the player’s inputs are mapped to a different part of the keyboard.

The second mapping illustrated in Input Map 1010 includes mapping of a screen area to a key. Specifically, a rectangle with corners at (100, 200) and (400, 500) is mapped to the F1 function key. The coordinates can be in image, window or screen pixels. This type of mapping is useful when a touch screen or pointing device is used as an input. This type of mapping can be used to map images within an overlay to specific game commands. This type of mapping can also be dependent on movement of a finger or cursor on a screen, or the number of points at which a touch screen is touched.

The third type of mapping illustrated in Input Map 1010 includes mapping of a controller button to a key. Specifically, a “Redbutton” is mapped to a “Spacebar.” This type of mapping can be used to map game controller inputs, such as those illustrated in FIGS. 9A-9C, to specific game commands.

The fourth type of mapping illustrated in Input Map 1010 includes mapping motion (“Right tilt”) of an input device to a key. This type of mapping can be used with joysticks, iPhones, iPads, the Nintendo Wii®, or other types of Clients 110 with motion or location sensing capabilities.

The fifth type of mapping illustrated in Input Map 1010 includes mapping of a movement of an object as detected by an external device, such as a camera, to a key “A.” This type of mapping is optionally used in Clients 110 that include radio or video motion detection systems. Using this type of mapping allows games designed to be played using a keyboard to be played using new types of input devices, such as the Kinect® from Microsoft.

A sixth type of mapping may be dynamic. For example, in embodiments wherein motion of a finger on a touch screen is mapped to joystick movements, where the finger first touches the screen may be dynamically mapped to the joystick “center” position. Screen positions just above, below, right, left, etc. of the first touch position are mapped to joystick positions up, down, right, left, etc., respectively. As such, a joystick movement can be entered by first placing a finger on the screen and then sliding the finger in some direction. Dynamic mapping may also include multi-finger touches to the screen. For example, a touch on a first point of the screen can determine the location and orientation of a mapping used to convert a second touch to a game command. Dynamic mapping is optionally facilitated by an overlay illustrating the mapping. For example, a semi-transparent joystick may be placed at the position the finger is first placed in the above example. In this case the second touch optionally represents movement of that joystick.

FIG. 11 illustrates methods of using an overlay, according to various embodiments of the invention. The overlay can be audio and or video. In the method illustrated, a video frame is generated using Graphics Processing Unit 235 and an overlay is added to this video frame prior to encoding. The methods illustrated by FIG. 11 are optionally performed in parallel using several of Game Logic 210, several of Video DEP 245 and/or several of Audio DEP 270.

Specifically, in a Generate Step 1110 a game environment is generated based on game rules and game commands received via Network 115 and Clients 110 from one or more game players. For example, a game player at Client 110A and a game player at Client 110B may each provide a game

36

command that is used to update a game state from which the game environment is generated. Generate Step 1110 may include Receive State Step 510 and Determine 3D Data Step 515, as discussed elsewhere herein. In some embodiments, Generate Step 1110 includes generating audio data based on events within the game, the game rules and/or received game commands.

In an optional Render Step 1120 the game environment generated in Generate Step 1110 is rendered using Graphics Processing Unit 235. Render Step 1120 is an embodiment of Render Step 520 discussed elsewhere herein. The result of Render Step 1120 is a video frame. Render Step 1120 is optional in embodiments where the result of Generate Step 1110 included only audio data and/or did not include a 3D game environment.

In a Provide Step 1130 the video frame generated in Render step 1120 and/or the audio data generated in Generate Step 1110 is provided to one or more encoding pipelines such as Video DEP 245A and/or Audio DEP 270A. For example, a video frame may be provided to Video DEP 245A while an audio frame is provided to Audio DEP 270A. Provide Step 1130 is an embodiment of Transfer Step 525.

In an Add Overlay Step 1140 a video overlay is added to the video frame provided in Provide Step 1130, and/or and audio overlay is added to audio data provided in Provide Step 1130. Add Overlay Step 1140 is typically performed using Overlay Logic 290 and/or Overlay Logic 295, as described elsewhere herein. More than one video overlay and more than one audio overlay are optionally added in Add Overlay Step 1140. The audio overlay optionally includes audio data received from one or more of Clients 110 as packets via Network 115. For example, the audio overlay can include voice data of game players received from two, three or more of Clients 110. The identities of the added overlay(s) are optionally dependent on a state of the video game.

In some embodiments Video Source 130 receives audio data from several, e.g., more than 3, 5 or 10 Clients 110 and Add Overlay Step 1140 includes determining which of the audio data should be used in a particular audio overlay. For example, while 10 game players may be using Video Source 130 only the voice data of two or three of these players may be included in the audio overlay. The determination of which audio data to include in an overlay is optionally based on a social relationship between game players. For example, one game player may select whose voice data he wishes to receive. The social relationships can include membership in a guild, party, being in a same battleground, or game instance, having avatars located near each other in the same game environment. In some embodiments, the volume of audio data is adjusted responsive to an in game distance between two avatars.

In an Encode Step 1150 the video frame and/or audio data, as well as any overlays are encoded. This is typically accomplished using Encoder 255 and/or Audio Encoder 280. Encode Step 1150 is an embodiment of Encode Step 535. Encode Step 1150 may be performed before or after Add Overlay Step 1140.

In a Deliver Step 1160 the video frame and/or audio data encoded in Encode Step 1150 is provided to one of Clients 110 via Network 115. The video frame and/or audio data are optionally provided to more than one location. For example, the same data can be provided to more than one game player. Deliver Step 1160 is an embodiment of Provide Step 545.

In an optional Map Step 1170 part of a video overlay applied in Add Overlay Step 1140 is mapped or remapped to a game command. For example, a region of the video frame to which the video overlay was added can be set in an Input Map

1010 as mapping to a specific game command. An overlay can be applied to more than one region of a display or a video frame. Each region is optionally mapped to a different game command. Input Map **1010** can be located on Video Server System **120** or one of Clients **110**.

In an optional Receive Step **1180** a mapped game command is received at or within Video Source **130**. The mapped game command is optionally a command that has been mapped from a region of a video frame to which an overlay was applied. The received game command may have been mapped at either one of Clients **110** or at Video Server System **120**. The received game command is optionally used to update a state of the video game.

FIG. **12** illustrates parts of Video Server System **120** including Sentinel Logic **1210**, according to various embodiments of the invention. In these embodiments undesirable input commands and input controls are blocked. Undesirable input commands and input controls include, for example, those whose operation expects that Game Logic **210A** is executing on Client **110A** rather than on Video Server System **120**. For example, an input command to select local video settings, manage local add-ons, save a game locally, manage audio settings, refresh a browser window, open a Universal Resource Locator in a browser window, etc. are potentially undesirable. These input commands can be redirected or otherwise blocked using the embodiments of Video Server System **120** illustrated in FIG. **12**. Input commands that should be redirected or otherwise blocked are referred to herein as “blocked input commands.” In some embodiments, the identities of blocked input commands are dependent on the state of a video game and/or on a game player’s account.

Some blocked input commands are associated with input controls included in the encoded video provided by Video Source **130** to Client **110A**. When an input command is blocked it is sometimes desirable to modify the associated input control such that a game player is not presented with an inoperable or improper control. In other words, if add-ons are not allowed then the game player should not be shown a “load add-ons” button. Input controls associated with blocked input commands can be grayed out, covered completely, changed to represent a different input command, or otherwise modified. Modified input controls are referred to herein as “blocked input controls.” Regarding both commands and input controls, the term “blocked” is meant to include generally both items that are completely eliminated as well as those that are merely redirected or changed. Blocked input controls are optionally animated or moving and can be moving controls in 2D, 3D or stereoscopic 3D.

Sentinel Logic **1210** is configured to detect the presence of blocked input controls. Detection is typically accomplished by monitoring the output of one or more elements of Video Server System **120** and comparing this output with a set of objects known to represent blocked input controls. This detection can occur at several points in the generation of encoded video (or still images). The set of objects is optionally stored in and retrieved from Storage **155**. In some embodiments the monitoring for blocked input controls occurs only every n th frame or image ($n > 1$). In some embodiments monitoring occurs only when the state of a game changes and/or a new image is generated. Sentinel Logic **1210** includes hardware, firmware and/or software stored on a computer readable medium.

In one example, detection occurs in a game environment generated by Game Logic **210**. This detection is accomplished by monitoring the generated game environments for text, shapes or textures known to represent blocked input controls. In another example, detection occurs at the output of

Graphics Processing Unit **235**, e.g., in memory of GPU **235** or in Shared Memory **240**. At this point, detection is accomplished by comparing the output with a bitmap known to represent blocked input controls. The bitmap may be compared with specific regions within the output or with all of the output. Thus, a blocked input command that could appear in a variety of alternative positions can be detected. Detection may, likewise, occur at any point within Video DEP **245A**. For example, Sentinel Logic **1210** can be configured to detect the presence of a blocked input control at the output of Overlay Logic **290**. Sentinel Logic **1210** is optionally further configured to detect the presence of a blocked input control in an overlay before the overlay is applied to an output of Graphics Processing Unit **235**. Sentinel Logic **1210** is optionally further configured to detect the presence of a blocked input control in video frames or images received from sources other than Graphics Processing Unit **235**.

The operation of Sentinel Logic **1210** may be responsive to optional State Logic **1220**. State Logic **1220** is configured to monitor the state of a video game (or some other computer application) and identify which input controls are blocked input controls based on the state. A specific input control may be a blocked input control in one state and a non-blocked input control in another state of the computer game. For example, a “save button” may be blocked when displayed on a top level menu and configured to initiate a command that would save the game locally, but unblocked when displayed on an avatar configuration menu and configured to initiate a command that would save an avatar configuration. Different sets of input controls can, thus, be blocked or non-blocked depending on the state of the video game. Sentinel Logic **1210** is configured to use the current set of blocked input controls when detecting the presence of blocked input controls. Further, the blocking of input controls and/or which input commands are blocked is optionally dependent on characteristics of a game player’s account, e.g., a premium account may be allowed controls and commands that a standard account is not.

The embodiments of Video Server System **120** illustrated in FIG. **12** optionally further include Command Processing Logic **1230**. Command Processing Logic **1230** is configured to redirect or otherwise block communication of blocked input commands to members of Game Logic **210**. In some instances, Command Processing Logic **1230** is also configured to redirect or otherwise block communication of blocked input commands to add-ons. As with Sentinel Logic **1210**, operation of Command Processing Logic **1230** is optionally controlled by State Logic **1220**. Thus, the set of input commands that are blocked by Command Processing Logic **1230** can be dependent on a video game state. Command Processing Logic **1230** includes hardware, firmware and/or software stored on a computer readable medium. Command Processing Logic **1230** is optionally disposed to examine input commands received by I/O Device **140** before they are communicated to parts of Video Source **130**.

Blocked input commands can be initiated by an input control included in the video or image output of Video Server System **120**, by keystrokes, by a browser add-on executing on Client **110A**, by game logic executing on Client **110A**, by a video game add-on, by a game controller, by a motion detection system, by a voice recognition system, or by any other input device. Blocked input commands can be blocked completely by Command Processing Logic **1230**. Alternatively, Command Processing Logic **1230** may be configured to map a blocked input command to a non-blocked input command or to redirect a blocked input command to a new destination. Input commands that are completely blocked are dropped

without further processing. Redirected input commands can be redirected to an add-on or to alternative game logic. For example, a “save” command directed at a part of Game Logic 210A configured to save a game on a client may be redirected to alternative game logic configured to save the game on a Video Server System 120. Input commands that are remapped are optionally remapped using a mapping similar to that illustrated in FIGS. 9A, 9B and 9C. Under the control of State Logic 1220, blocked input commands can be completely blocked in a first game state, redirected to a first location in a second game state, redirected to a second location in a second game state, remapped to a first alternative input command in a third state, remapped to a second alternative input command in a fourth state, or any combination of the above.

In some embodiments, a blocked input control is detected by Sentinel Logic 1210 in the output of GPU 235 within Shared Memory 240. The blocked input control includes a “load Add-on” button and a data entry field configured for a game player to enter a local file name and path of an add-on. Overlay Logic 290 is used to modify the “Load Add-on” button to read “Load Remote Add-on” and the data entry field is covered using a background colored bitmap. When the “Load Remote Add-on” button is clicked by a game player a blocked “Load Add-on” input command is generated at Client 110A and sent to Video Server System 120. This input command can be communicated as a text stream, as a screen coordinate, as a binary object, and/or the like. At Video Server System 120 the “Load Add-on” input command is redirected by Command Processing Logic 1230 to game logic configured to load one of a set of pre-approved add-ons on Video Server System 120. This game logic is optionally itself an add-on to Game Logic 210A. In some embodiments, an audio overlay including for example a “warning” sound is added to the audio stream at the same time that a video overlay is added to a video frame or when a blocked input control is activated by a game player.

In some embodiments, some blocked input commands are configured to control the operation system or some other software application of Client 110A. In these embodiments the blocked input command is redirected to optional Client Control Logic 1240. Client Control Logic 1240 includes hardware, firmware, and/or software stored on a computer readable medium, and can be disposed on Video Server System 120 and/or one of Clients 110. In some embodiments Client Control Logic 1240 includes a browser add-on configured to be added to a browser of Clients 110.

In some embodiments, Client Control Logic 1240 is configured to open and/or close windows, e.g., browser windows (or tabs) or other window types, on Client 110A. These windows are optionally provided with a network address or universal resource locator configured to retrieve images or video from an instance of Video Server System 120. As such, Client Control Logic 1240 can be used to open a new window on Client 110A and automatically provide a website and/or streaming video to that window. The opened window is optionally configured for logging into a social networking website or configured for local execution of a version of Game Logic 210A.

In one example, Command Processing Logic 1230 is configured to redirect a blocked input command to open a Windows Explorer® window on Client 110A such that a user can select a local file. This blocked input command is redirected to Client Control Logic 1240 (optionally among other places) which opens a browser window on Client 110A rather than a Windows Explorer window. The browser window is provided with a URL to a video output of Video Server System 120, e.g. a port of I/O Device 145, at which a video of a virtual file

system of Video Server System 120 is provided. The user (or game player) can navigate the virtual file system to select a file by clicking on input controls within the provided video. Either Client Control Logic 1240 or Command Processing Logic 1230 may be configured to initiate generation of the video of the virtual file system using video Server system 120. A result of this process is that a second browser window is opened and the images (e.g., video) provided to that window gives the same look and feel as if the video game or other software application were executed locally on Client 110A.

Video Server System 120 optionally further includes Discovery Logic 1250 configured to identify blocked input controls within graphics output of Video Source 130. Once identified these blocked input controls can be modified in future output of Video Source 130 using the systems and methods discussed herein. The operation of Discovery Logic 1250 is optionally automatic. Discovery Logic 1250 includes hardware, firmware, and/or software stored on a computer readable medium.

Discovery Logic 1250 may be configured to identify blocked input controls using a variety of different approaches. In some embodiments, Discovery Logic 1250 compares input commands received from Client 110A with a list of undesirable input commands. When a match is found the blocked input control on video already provided to Client 110A is identified, for example, by matching a bitmap, optical character recognition, or use of the coordinates of a game player’s click on the video. In some cases the full identification of a blocked input control requires the identification of multiple click locations.

In some embodiments, Discovery Logic 1250 monitors activities of Game Logic 210A in order to determine which input commands are blocked input commands. For example, Discovery Logic 1250 may watch for calls made by Game Logic 210A to a file system or to video hardware, and observe which input commands were recently received by Game Logic 210A that may have caused these calls. Once the blocked input commands are identified, the associated blocked input control can be identified as described elsewhere herein. In some embodiments, Discovery Logic 1250 is configured to examine the output of add-ons.

FIG. 13 illustrates methods of generating game video, according to various embodiments of the invention. These methods are optionally performed using elements of the embodiments illustrated in FIG. 12. The generated game video includes an overlay configured to modify a blocked input control.

In Provision DEP Step 450, an encoding pipeline such as Video DEP 245A is provisioned. Provision DEP Step 450 is optional in embodiments wherein an encoding pipeline, including overlay logic, is already available. Provision DEP Step 450 may include other steps illustrated in FIG. 4. Further details of Provision DEP Step 450 are described elsewhere herein.

As described elsewhere herein, Generate Step 1110 includes generating a game environment based on game rules and input commands received from one or more of Clients 110. The game environment is rendered to video frames using Graphics Processing Unit 235 in Render Step 1120 and these video frames are provided to one or more of Clients 110 in Provide Step 1130.

In a Monitor Step 1315 the video frames are monitored for the presence of blocked input controls. This monitoring is optionally performed using Sentinel Logic 1210. The monitoring can occur prior to the inputs to Graphics Processing Unit 235, e.g., by monitoring the game environments. Alternatively, or additionally, the monitoring can occur by exam-

41

ining video frames generated by Graphics Processing Unit 235. The frames can be monitored at memory internal to Graphics Processing Unit 235, at the outputs of Graphics Processing Unit 235, within Shared Memory 240, within Video DEP 245A, and/or the like. For example, monitoring may occur at Shared Memory 240 and also after the inputs to Overlay Logic 290. In some embodiments only every n^{th} ($n > 1$) frame is monitored. Frames are optionally monitored by searching for image characteristics, features and/or bitmaps known to represent blocked input controls and, thus, indicate a need to block input controls, certain keystrokes or other inputs when the frame is displayed. As described elsewhere herein, State Logic 1220 is optionally used to control the monitoring based on a game state. In various embodiments, using the game state allows the number of bitmaps that are monitored for at a time to be reduced by at least 50, 75, or 90%. Monitor Step 1315 is optionally applied to video frames or still images that are generated or retrieved through means other than Graphics Processing Unit 235. For example, Monitor Step 1315 can be applied to two dimensional images, video retrieved from storage, live video generated using a camera, video received from a member of Clients 110, video received from an external source, and/or the like.

In optional Add Overlay Step 1140 an overlay is applied to an image (e.g., video frame) using Overlay Logic 290. As described elsewhere herein, the overlay can include input controls. The overlay is optionally therefore also monitored (in Monitor Step 1315) for the presence of blocked input controls. The overlay can be monitored before or after it is applied to the image.

In Encode Video Step 1150 and Deliver Step 1160 video frames are encoded using one or more of Encoders 255 and I/O Device 145. The video frames are eventually received for display on one or more of Clients 110. These steps are described further elsewhere herein.

In a Block/Redirect Step 1320 the communication of a blocked command received from a member of clients is blocked or redirected and, thus, prevented from reaching a member of Game Logic 210. The blocked command is optionally redirected to an add-on, alternative game code, and/or the like. The blocked command is optionally then be replaced by an allowed, alternative command. In some embodiments, the blocked command is redirected to Discovery Logic 1250 configured to attempt to identify the blocked input control. When this occurs, the Discovery Logic 1250 optionally attempts to identify the blocked input control on the graphic output that caused generation of the blocked command. If identified, the blocked input control can, thus, be blocked in the future using techniques described herein.

In Map Step 1170 part of a video overlay applied in Add Overlay Step 1140 is mapped or remapped to an alternative game command or no game command at all. Operation of the covered blocked input control is typically changed such that it does not generate a blocked command. Further details of this step are described elsewhere herein. The identification and redirection of a blocked input control in Steps 1320 and 1170 is optionally performed automatically. The blocked input control can be a default part of a game or the result of an add-on.

In an optional Control Window Step 1325, a browser window or browser tab is opened or closed on Client 110A using Client Control Logic 1240. Control Window Step 1325 is optionally performed in response to the blocked or remapped input command received in Block/Redirect Step 1320. Block/Redirect Step 1320, Map Step 1170 and Control Window Step 1325 are optionally performed independently of other steps illustrated in FIG. 13. In some examples, the opened browser

42

window is directed to a help web page, a registration or purchase web page, an add-on web page, an account management web page, a game selection web page, a user or add-on specified web page, an alternative video stream including game video, and/or the like.

FIG. 14 illustrates methods of identifying blocked input controls, according to various embodiments of the invention. In these methods a region of a game video frame is identified as including a blocked input control. The blocked input control can then be modified using an overlay in future video frames and/or the command generated by the blocked input control can be redirected or otherwise blocked.

In an Identify Step 1410, a blocked input command (a control command to be blocked) is identified in commands received from a member of Clients 110. For example, in various embodiments, the command is identified prior to or after being communicated to a member of Game Logic 210. Identification prior to receipt by Game Logic 210 typically requires that the command be recognized as one that could cause an undesirable result. For example, a command known to be for accessing video settings. In some cases commands can only be identified as being blocked input commands based on what Game Logic 210 does in response to receiving the command. For example, if Game Logic 210A performs an undesirable file access or attempts to modify a video driver after a command is received, then that command may automatically be considered a blocked input command. Identify Step 1410 is optionally performed automatically by Discovery Logic 1250. As described elsewhere herein, whether or not an input command is a blocked input command is optionally state dependent.

In a Find Step 1415 a region (e.g., frame coordinates) within a video frame that resulted in the blocked input command identified in Identify Step 1410 is found. This is optionally accomplished by identifying the location of a user "click" on the frame and/or by recognizing specific objects, such as buttons or icons within the frame. Find Step 1415 is optionally performed automatically using Command Processing Logic 1230, Sentinel Logic 1210 and/or Discovery Logic 1250.

In an optional Select Overlay Step 1420 an overlay is selected to modify the region of future video frames found in Find Step 1415. The overlay may modify all or part of a blocked input control depending on how much of the input control has been found. More than one occurrence of Identify Step 1410 and Find Step 1415 may be required to fully identify a blocked input control. The overlay is optionally selected from a group of standard overlays, from a "grayed out" version of the blocked input control, or from an image dynamically generated based on contents of the video frame. The selected overlay is configured to be applied to future instances of the video frame as monitored by Sentinel Logic 1210.

In an optional Store Select Overlay Step 1425 the selected overlay is stored in a location accessible to Overlay Logic 290. Store Select Overlay Step 1425 is optionally in embodiments wherein the selected overlay is already stored in a location accessible to Overlay Logic 290.

In an optional Determine Mapping Step 1430 a mapping is determined from the modified input command control to an input command. For example, a blocked command control that resulted in sending of a video settings command may be remapped to alternative logic configured to offer allowed video operations. The blocked command control may or may not be modified by an overlay. Remapping may be accomplished using Command Processing Logic 1230 and/or Discovery Logic 1250.

43

In an optional Store Mapping Step **1435** the mapping determined in Determine Mapping Step **1430** is stored in a location accessible to Command Processing Logic **1230**.

Several embodiments are specifically illustrated and/or described herein. However, it will be appreciated that modifications and variations are covered by the above teachings and within the scope of the appended claims without departing from the spirit and intended scope thereof. For example, while the examples provided herein are directed to video games, the systems and methods discussed can be applied to other interactive video systems and other computer applications. For example, Game System **100** may be adapted to be a generalized application system including application logic in place of Game Logic **210A** and/or **210B**. This application logic can include accounting applications, computer aided design applications, word processing applications, video on demand applications, drawing applications, image editing or viewing applications, browsing applications, communication applications, spreadsheet applications, database applications, presentation management applications, and/or the like. While the examples provided herein are primarily directed at the delivery of video, the systems and method described herein can be applied to management of audio commands, audio overlays, the generation of audio frames and delivery of audio streams. In some embodiments, audio and video streams are generated and delivered together as part of an audio-video stream.

The embodiments discussed herein are illustrative of the present invention. As these embodiments of the present invention are described with reference to illustrations, various modifications or adaptations of the methods and or specific structures described may become apparent to those skilled in the art. All such modifications, adaptations, or variations that rely upon the teachings of the present invention, and through which these teachings have advanced the art, are considered to be within the spirit and scope of the present invention. Hence, these descriptions and drawings should not be considered in a limiting sense, as it is understood that the present invention is in no way limited to only the embodiments illustrated.

Computing systems referred to herein can comprise an integrated circuit, a microprocessor, a personal computer, a server, a distributed computing system, a communication device, a network device, or the like, and various combinations of the same. A computing system may also comprise volatile and/or non-volatile memory such as random access memory (RAM), dynamic random access memory (DRAM), static random access memory (SRAM), magnetic media, optical media, nano-media, a hard drive, a compact disk, a digital versatile disc (DVD), and/or other devices configured for storing analog or digital information, such as in a database. The various examples of logic noted above can comprise hardware, firmware, or software stored on a computer readable medium, or combinations thereof. A computer-readable medium, as used herein, expressly excludes paper. Computer-implemented steps of the methods noted herein can comprise a set of instructions stored on a computer-readable medium that when executed cause the computing system to perform the steps. A computing system programmed to perform particular functions pursuant to instructions from program software is a special purpose computing system for performing those particular functions. Data that is manipulated by a special purpose computing system while performing those particular functions is at least electronically saved in buffers of the computing system, physically changing the special purpose computing system from one state to the next with each change to the stored data.

44

The logic discussed herein may include hardware, firmware and/or software stored on a computer readable medium. This logic may be implemented in an electronic device to produce a special purpose computing system.

What is claimed is:

1. A video server system comprising:

an application engine configured to execute a first computer application;

a graphics processing unit configured to generate a graphic output in response to graphics data received from the first computer application;

sentinel logic configured to monitor the graphic output and determine that a blocked input control is included within the graphic output, the sentinel logic configured to send a blocked command to disable the blocked input control, the sentinel logic provides instruction to adjust the graphic output so as to modify the blocked input control, wherein the sentinel logic is configured to determine that the blocked input control is included within the graphic output by identifying a bit pattern that is predefined for the blocked input control;

an overlay logic that receives instruction from the sentinel logic to adjust the graphic output, the overlay logic is configured to process the modifying by placing an overlay on a portion of the graphic output that includes the blocked input control, the overlay configured to change an appearance of the blocked input control, wherein the sentinel logic is further configured to examine the overlay to ensure absence of any blocked input control in the overlay prior to the overlay being placed on the portion of the graphic output that includes the blocked input control;

an encoding pipeline configured to generate an encoded video stream using the graphic output;

an output logic configured to communicate the encoded video stream to a remote client via a communication network, the encoded video stream including the graphic output that includes the modified blocked input control; and

a microprocessor configured to execute the sentinel logic at the video server system.

2. The system of claim 1, wherein the first computer application is a video game.

3. The system of claim 1, wherein the graphic data includes a three dimensional game environment and the graphic processing unit is configured to render the graphic output based on the game environment and a point of view.

4. The system of claim 1, wherein the sentinel logic is configured to search for the bit pattern by examining a plurality of bit patterns within the graphic output, and to search for different members of the plurality of bit patterns depending on a state of the first computer application.

5. The system of claim 1, wherein the sentinel logic is configured to monitor only a fraction of graphic output frames generated by the graphics processing unit in response to the first computer application.

6. The system of claim 1, wherein the sentinel logic is configured to notify the overlay logic when a blocked input control is determined to be present in the graphic output.

7. The system of claim 1, wherein the sentinel logic is configured to determine that a blocked input control is included within the graphic output by monitoring memory within the graphics processing unit or by monitoring memory shared by more than one encoding pipeline.

8. The system of claim 1, further comprising state logic configured to change operation of the sentinel logic respon-

45

sive to a state of the first computer application, the change including selecting which input controls are blocked input controls.

9. The system of claim 1, further comprising discovery logic configured to identify previously unidentified blocked input controls within the graphic output. 5

10. The system of claim 9, wherein the discovery logic is configured to identify the previously unidentified blocked input controls by monitoring actions of the first computer application. 10

11. The video server system of claim 1, wherein the sentinel logic is configured to send the blocked command to the first computer application so that subsequent graphics data received from the first computer application does not include the blocked input control. 15

12. A video server system comprising:

an application engine configured to execute a first computer application;

a graphics processing unit configured to generate a graphic output in response to graphics data received from the first computer application; 20

sentinel logic configured to monitor the graphic output and determine that a blocked input control is included within the graphic output, the sentinel logic configured to send a blocked command to disable the blocked input control, the sentinel logic provides instruction to adjust the graphic output so as to modify the blocked input control, wherein the sentinel logic is configured to determine that the blocked input control is included within the graphic output by identifying a bit pattern that is predefined for the blocked input control; 25

an overlay logic that receives instruction from the sentinel logic to adjust the graphic output, the overlay logic is 30

46

configured to process the modifying by placing an overlay on a portion of the graphic output that includes the blocked input control, the overlay configured to change an appearance of the blocked input control, wherein the sentinel logic is further configured to examine the overlay to ensure absence of any blocked input control in the overlay prior to the overlay being placed on the portion of the graphic output that includes the blocked input control;

an encoding pipeline configured to generate an encoded video stream using the adjusted graphic output;

an output logic configured to communicate the encoded video stream to a remote client via a communication network, the encoded video stream including the graphic output that includes the modified blocked input control; 15

client control logic configured to open a browser window or browser tab on the remote client in response to receiving a blocked command from the remote client; and

a microprocessor configured to execute the sentinel logic.

13. The system of claim 12, further comprising command processing logic configured to communicate the blocked command received from the sentinel logic to the first computer application. 25

14. The system of claim 12, wherein the browser window or browser tab is configured for local execution of the first computer application on the client.

15. The system of claim 12, wherein the client control logic is further configured to load a website into the browser window or browser tab. 30

16. The system of claim 12, wherein the client control logic is further configured to log into a social networking website.

* * * * *